




SEMIANNUAL REPORT

October 1, 1996 through March 31, 1997



*NASA Contract No. NAS1-19480
May 1997*

*Institute for Computer Applications in Science and Engineering
NASA Langley Research Center
Hampton, VA 23681-0001*

Operated by Universities Space Research Association



*National Aeronautics and
Space Administration*

*Langley Research Center
Hampton, Virginia 23681-0001*

DTIC QUALITY INSPECTED 3

19970701 059

CONTENTS

Page

Introductionii

Research in Progress

Applied and Numerical Mathematics 1

Fluid Mechanics 14

Applied Computer Science 27

Reports and Abstracts 41

ICASE Interim Reports 45

ICASE Colloquia 46

Other Activities 48

ICASE Staff 49

DTIC QUALITY INSPECTED 3

INTRODUCTION

The Institute for Computer Applications in Science and Engineering (ICASE)* is operated at the Langley Research Center (LaRC) of NASA by the Universities Space Research Association (USRA) under a contract with the Center. USRA is a nonprofit consortium of major U. S. colleges and universities.

The Institute conducts unclassified basic research in applied mathematics, numerical analysis and algorithm development, fluid mechanics, and computer science in order to extend and improve problem-solving capabilities in science and engineering, particularly in the areas of aeronautics and space research.

ICASE has a small permanent staff. Research is conducted primarily by visiting scientists from universities and industry who have resident appointments for limited periods of time as well as by visiting and resident consultants. Members of NASA's research staff may also be residents at ICASE for limited periods.

The major categories of the current ICASE research program are:

- Applied and numerical mathematics, including multidisciplinary design optimization;
- Theoretical and computational research in fluid mechanics in selected areas of interest to LaRC, such as transition, turbulence, flow control, and acoustics;
- Applied computer science: system software, systems engineering, and parallel algorithms.

ICASE reports are considered to be primarily preprints of manuscripts that have been submitted to appropriate research journals or that are to appear in conference proceedings. A list of these reports for the period October 1, 1996 through March 31, 1997 is given in the Reports and Abstracts section which follows a brief description of the research in progress.

*ICASE is operated at NASA Langley Research Center, Hampton, VA, under the National Aeronautics and Space Administration, NASA Contract No. NAS1-19480. Financial support was provided by NASA Contract Nos. NAS1-19480, NAS1-18605, NAS1-18107, NAS1-17070, NAS1-17130, NAS1-15810, NAS1-16394, NAS1-14101, and NAS1-14472.

RESEARCH IN PROGRESS

APPLIED AND NUMERICAL MATHEMATICS

BRIAN G. ALLAN

Base flow design controlled Navier-Stokes system

A renewed interest in the area of active flow control has been sparked by new technological advancements in micro actuators and sensors. These micro sensors and actuators have shown a promise for creating advanced fluid dynamical systems by incorporating feedback control. Of particular interest to NASA is the delay of separation in high-lift systems. Before doing experimental tests, one would like to know where to place these micro devices in order to minimize the control effort. In order to gain insight into the areas where sensors and actuators should be placed, a reduced order model of the Navier-Stokes equations will be used in an LQR optimal feedback control design. By assuming sensing and actuation everywhere, a spatially varying feedback gain can be calculated using optimal control theory. This feedback gain will help guide the placement of sensors and actuators on a high-lift system by identifying the spatial regions where feedback sensitivity and actuation are large.

Developing a model to be used in the control design is a key element to this investigation. By linearizing the Navier-Stokes equations about some desired steady flow state, a set of equations describing the stability of the desired flow can be found. These equations are then discretized to obtain a finite dimensional model of the system. It is important that we linearize about a steady state solution so that the steady control input required to maintain the desired flow field is minimized. Since the flow state we desire is unstable, a steady state solution to the Navier-Stokes equations can not be obtained through time evolution. Currently, a code is being developed which will use an inexact Newton-Krylov method to find the desired unstable steady state solution. The first estimate to the steady state is developed by matching an outer potential flow solution to a boundary layer profile. This initial guess is used as a starting point for the Newton-Krylov iterations to a steady state solution.

Future work will begin by incorporating the designed base flow into the linear model currently being developed by J. Loncaric at ICASE. Once the model has been developed, it will be used to provide insight for the placement of sensors and actuators by using optimal control theory.

This work is done in collaboration with J. Loncaric at ICASE.

EYAL ARIAN

Coupling theory for multidisciplinary analysis and optimization

The objective of this work is to develop quantitative estimates of how tightly coupled is a given system of coupled PDEs with respect to analysis, sensitivity, and optimization computations. This work is relevant to algorithm development for problems in multidisciplinary analysis and optimization.

The approach is to derive expressions for the coupling based on the partial differential operators involved and then to estimate an upper bound for these expressions. The predictions of the analysis were verified numerically on simplified model problems including a non-linear algebraic system and a system of linear PDEs (static potential flow over a plate).

Future plans are to further investigate the applications of the above methodology to MDO problems.

Optimization problems for radome design

The objective of this work is to determine the optimal thickness and permittivity of a multilayer Radome such that some desired properties are attained.

Our approach is to pose these problems as optimization problems with possible inequality constraints on the design variables (thickness and permittivity) and solve them with gradient based algorithms. We have been able so far to determine the optimal thickness of a symmetric three layer Radome, averaging over a specific frequency band (around 94 GHz) and incident angles, such that the transmission of radiation through the Radome is maximal.

Future plans are to solve inverse problems to determine Radome parameters given measurements of the transmission, reflection data, and layer thickness and also to work on anisotropic problems.

This work is done in collaboration with F.B. Beck, C.R. Cockrell, C.J. Reddy and M.D. Deshpande.

Acceleration techniques for aerodynamic optimization

Our goal is to develop efficient, simple to use, algorithms for shape optimization of airfoils.

Our approach is based on the following principles:

- Use pointwise (infinite dimensional) representation of the airfoil.
- Derive the adjoint equations in the continuous level and then discretize (to compute the sensitivity gradients).
- Precondition the gradient direction by approximating the inverse of the Hessian in the continuum level.
- Estimate the step size by the above approximation thus avoid a line search.

We have implemented successfully this new approach on optimal shape design of airfoils under subsonic, transonic, and supersonic flow speeds using Euler equations (with TLNS3D). We find the new method computationally efficient and simple to apply. We intend to further develop this method.

This work is done in collaboration with V. Vatsa of NASA Langley.

H.T. BANKS

Application of smart material technology to damage detection in plate-like structures

The focus of this research is to apply smart material technology to the non-destructive detection and characterization of damage in plate-like structures using vibration data from self-sensing, self-actuating structures. Among the different types of smart materials currently studied, structures

with bonded piezoelectric ceramic patches are of particular interest. When an electric field is applied, piezoelectric patches induce strains in the materials they are bonded to and, inversely, they produce a voltage when a deformation occurs in the material. As a consequence, these patches can act both as actuators and sensors, providing the host structure with smart material capabilities.

We have developed a mathematical model for a damaged 2-D plate-like structure with piezoelectric patches bonded onto its surface. Computational methods for the detection of damage and characterization of physical parameters such as stiffness, damping, or mass density are being developed.

We expect to conduct some experiments to characterize damaged structures and verify our computational algorithms with experimental data.

This research is conducted in collaboration with Pierre Emeric (North Carolina State University) and W.P. Winfree (NASA Langley Research Center).

DAVID GOTTLIEB

Computational methods in material sciences

The research in material sciences considers optimization for a set of two equations which describe chemical vapor infiltration (CVI). In this process a vapor-phase precursor is transported into the porous preform, and a combination of gas and surface reactions leads to the deposition of the solid matrix phase.

We derived a simple set of two equations that models the infiltration processes, and we showed how to get these equations as a subcase of the general system. These are nonlinear partial differential equations. We also derived initial and boundary conditions for the model equations. We also presented a detailed mathematical analysis concerning the behavior of the solutions in space and time. This analysis provides insights into the behavior of the process. We define the concept of a successful process and get conditions on the parameters of the problem for a process to be successful. In particular we formulated an optimization problem for the minimum time it takes for the process to settle. Numerical experiments are carried out to validate the theory. We also instigated how to design the experimental parameters to get faster successful processes.

Future work will involve more sophisticated models in general geometries.

JAN S. HESTHAVEN

High order multi-domain methods for wave problems in complex geometries

Traditional high-order/spectral methods suffer from having a fixed distribution of grid points, thereby making it difficult to apply such methods for solving problems with strong internal layers or problems in complex geometries. However, splitting the computational domain into several geometrically simple domains has proven to be a powerful way of overcoming these restrictions on the applicability of high-order/spectral methods in order to solve problems of interest to industry. In addition, such multi-domain methods lend themselves to efficient implementation on contemporary parallel computers. The objective of this research has been to identify and solve the problems associated with such an approach with the emphasis being on the solution of wave dominated problems such as the compressible Navier-Stokes equations and Maxwell's equations.

One of the most important issues to address, when considering the application of high-order methods, is the development of consistent and stable ways of dealing with complex boundary conditions. We have found that imposing the boundary conditions as well as interface conditions in between subdomain in a weak form, using a penalty term, has several attractive features, e.g., the full semi-discrete schemes may be proven asymptotically stable and very general complex boundary operators can be dealt with in a very simple manner. Using this approach we have developed and implemented multi-domain methods, based on hexahedrals, for the solution of the compressible Navier-Stokes equations as well as Maxwell's equations. Moreover, we have developed high-order stable methods for the solution of hyperbolic problems on triangular domains using the penalty approach, thereby establishing a new building block for the construction of high order unstructured grid methods. Recently, we have introduced wavelet analysis, contributing as a local error estimator, into the multi-domain framework to arrive at p -adaptive methods for the accurate and efficient solution of wave dominated problems.

In the near future, we aim at refining the wavelet analysis for the present purpose and to implement the developed schemes for the solution of three-dimensional unsteady problems, possibly on a parallel computer, to address problems in gas dynamics, electromagnetics and, as an extension, aero-acoustics.

This research was completed as a joint effort with L. Jameson (ICASE) and D. Gottlieb (Brown University).

ANGELO IOLLO

Study of wave propagation in coupled aeroelastic systems

The noise emission inside a fuselage is known to be linked to the presence of rib stiffeners in the panel structure. Perturbations of an aerodynamic nature can excite panel vibrations. The vibrations are then scattered by the stiffeners. Previous analysis have studied the coupling between the linear beam equation and the wave equation, showing that the noise level is a certain function of the ratio between the wave speed in the beam and the wave speed in the fluid. The presence of a mean flow over the fuselage was ignored in previous work and is the subject of this study.

We first considered a quasi-one-dimensional aeroelastic system consisting of the quasi-one-dimensional Euler equations and the linear beam equation. We linearized the system and solved the dispersion relation. The analytical results obtained were confirming by numerical simulations using the same flow-beam model. Subsequently, we studied a two-dimensional system in which the fluid was governed by the linear compressible potential equation. A Fourier analysis of the system, and some approximations, yielded the analytical dependence on Mach number of the production of noise by scattering at rib stiffeners.

A numerical simulation showing some relevant cases, studied analytically, will complete the study.

This research was undertaken jointly with Manuel D. Salas.

LELAND M. JAMESON

A wavelet optimized, adaptive multidomain method

Error control, grid adaptation, and order selection can all be accomplished quite naturally by using wavelet analysis. That is, wavelet analysis yields a set of coefficients indicating energy levels of flow variables as a function of scale and location throughout the domain. Such information is ideal for grid selection. Furthermore, wavelets analysis can distinguish between large smooth structures and small rough structures throughout domain, and such information is ideal for selecting the order of accuracy of differencing operators. Finally, by combining grid generation and order selection, one can keep the maximum error down to any prescribed user-selected tolerance in an efficient manner.

An adaptive method has been created which uses wavelets to optimally choose the numerical grid point density and the order of the numerical approximation. In addition, the method is designed to work efficiently in a parallel environment. The method breaks the domain into subdomains, and on each subdomain the optimal grid density and the optimal order is fixed. The grid density and order are selected not only for optimal propagation, but also for automatic load balancing among the subdomains allowing for efficient parallel implementation. This method will work for arbitrary geometries by mapping each of the domains.

The next step will be to extend the method to higher dimensions and to modify for complicated geometries.

This research was conducted in collaboration with J.S. Hesthaven (Brown University).

MICHAEL LEWIS

Nonlinear programming for engineering design

Great progress has been made in the ability to accurately and faithfully simulate the behavior of physical and engineering systems. However, the enormous computational cost of such simulations makes it impractical to rely exclusively on high-fidelity simulations for the purpose of design optimization. Our objective is to make as much use as possible of models of lower physical fidelity but lower computational cost, with only occasional recourse to expensive, high-fidelity simulations. This is in keeping with one tenet of nonlinear programming: that one should try to avoid doing too much work when far from an optimizer. Moreover, the use of approximation models is in keeping with engineering practice, where models of lower fidelity are widely used in preliminary design to explore the design space. Furthermore, approximations based on experimental data enable us to include non-computational information in the optimization process.

We have developed a trust region framework for the use of general non-quadratic approximations in optimization that insures robust global behavior. The significance of this work is that it is one of the few systematic approaches to the use of non-quadratic approximations and surrogates in nonlinear programming and the first attempt to provide an analytical justification for such strategies. We are currently developing software to implement specific classes of approximations as well, in particular, software for universal kriging and the *beta*-correlation method of Chang, Haftka, Giles, and Kao.

In addition to this software to construct and use approximations and surrogates, we have continued the development of basic nonlinear programming algorithms. We have extended pattern

search methods to linearly constrained optimization and continue work on our large-scale trust-region barrier method code. The accomplishment here is that we continue to extend the variety of practical nonlinear programming techniques available for use in engineering design.

The next step is to apply the approximation management approach to realistic applications. We have identified several structural optimization problems as suitable test cases. We also plan to test the validity of the universal kriging model on wind tunnel data.

This research was done, in part, in collaboration with Natalia Alexandrov (NASA LaRC) and Virginia Torczon (William & Mary).

JOSIP LONČARIĆ

Spatial structure of separation feedback control

Flow separation over wings at high angles of attack limits lift, which is particularly important at takeoffs and landings. Low limit on lift necessitates larger and heavier wings, which has a disproportionately large effect on the useful aircraft payload. Several techniques (heating, blowing, sound) are known to delay flow separation, with open-loop unsteady inputs being more effective than steady inputs. New technologies (e.g. MEMS) hold the promise of constructing smart aircraft systems capable of using feedback control to maintain the desired flow state. While much *can* be done to the flow, the question of what *should* be done remains open. Traditionally, control design is carried out after the system is already specified. We intend to compute optimal feedback gains not with the goal of implementing them, but with the goal of providing new insight helpful in designing smart high lift systems.

This project will compute the minimum effort distributed flow control feedback operator and investigate its spatial structure. Those spatial regions where feedback sensitivity and action are large will guide placement of a finite number of sensors and actuators, suitable for control implementation later. Our approach begins with the recent work on optimal feedback control of incompressible Navier-Stokes equations. The main remaining obstacle is computational cost, which we aim to reduce by using spectral flow representation on an efficient grid. A reduced system model which extracts the significant dynamics will be constructed using Krylov subspace methods and computational techniques borrowed from stability analysis. This approach is equivalent to the Hankel norm optimal reduced model since control and sensing operators are initially identities. The laminar 2D exterior flow problem has been reformulated in terms of a new state variable which simplifies the application of exact artificial boundary conditions (replacing the boundary condition at infinity). This representation also simplifies the interpretation of the implicit vortex sheet generated at the solid boundary, a process potentially responsible for flow instabilities even if the perturbation growth through advection of the nominal flow vorticity is neglected. Numerical properties of the discretized operators involved are under investigation. The nearly steady but unstable nominal flow state is being constructed (in collaboration with Brian G. Allan of ICASE) by Newton-Krylov methods using as the initial guess the exterior potential flow with an approximately steady boundary layer.

The problem is conceptually defined and the numerical algorithm is largely worked out. We intend to complete the investigation of the numerical properties of the chosen discretization. Preliminary runs will be done at coarse resolution (32^2 fully dealiased), and the spatial structure of these

feedback operators will be used to tune the scheme for the final runs (up to 256^2 fully dealiased). These distributed control results will enable us to guide further work in boundary control. This numerical study is also expected to lead to useful analytical approximations in the feedback control of boundary layers.

Placement of sensors and actuators in noise control

Interior noise on aircraft is largely produced by vibration of fuselage panels in response to forcing by exterior pressure fluctuations. Passive damping of such vibration exacts a significant weight penalty, which can be alleviated by active control. Successful active noise control strategies have been demonstrated by attaching piezoelectric patches to the panel surface, but the question of optimal placement of sensors and actuators remains unanswered. Motivated by this problem, we aim to develop a model-based optimal control theory approach to the design of such distributed parameter systems.

Structural vibration of panels and interior acoustic radiation can be modeled by differential operators which belong to the class of spectral systems. If one could measure everything and act everywhere, what should one do? In this thought experiment the control input is allowed to act everywhere, which allows simultaneous diagonalization of both system dynamics and control. The infinite dimensional LQR problem therefore decouples into an infinite series of low dimensional LQR problems which are easy to solve, each of which contributes a term to the optimal feedback operator K . Spatial regions where K exhibits large sensitivity or action suggest good locations for sensors and actuators. We have obtained analytic expressions for K (in the distribution sense) for a simply supported plate and several control objectives. The results indicate that positional stabilization requires control of only the first few modes, but that minimizing the plate velocity, acceleration or energy requires high bandwidth. Furthermore, high frequency behavior of the optimal feedback gains at low internal damping is a singular perturbation of the undamped case. Human ears are sensitive only to certain range of frequencies which means that the required control bandwidth is also limited. The spatial structure of optimal K corresponding to a bandwidth limited objective function has been obtained analytically for a freely supported beam, indicating highly localized behavior. Without bandwidth limits, for several objectives tested the optimal K exhibits singularities along the line representing collocated sensors and actuators.

We intend to extend these results to the clamped plate case through numerical experiments. We also intend to explore coupling of structural and acoustic modes, as well as different input disturbances.

DIMITRI J. MAVRIPLIS

Adaptive meshing for mixed element unstructured meshes

The use of hexahedral and prismatic elements in unstructured meshes in addition to traditional tetrahedral elements offers the possibility of increased accuracy and reduced overheads. In previous work, a discretization and solution strategy for mixed-element unstructured meshes was developed, and gains in efficiency due to the use of different element types were demonstrated. The present work seeks to develop adaptive meshing techniques for mixed element unstructured meshes. The ability to adaptively refine mixed element meshes is important since adaptive meshing represents one of the principal advantages of unstructured meshes.

Mesh refinement is achieved by element subdivision. For each element type, the permissible subdivision patterns (isotropic and anisotropic subdivision) are defined, classified and incorporated into a generic element subdivision library. The hierarchical subdivision history is also stored. This enables de-refinement by retracing the subdivision history. More importantly, the hierarchical information is utilized to avoid multiple levels of anisotropic refinements. For anisotropically refined cells (which occur at interfaces between refined and unrefined regions) any further refinement requires the deletion of the current anisotropic children cells, and the isotropic refinement of the corresponding parent cells. This is necessary to avoid degradation of mesh element quality with additional refinement levels. Adaptive refinement of fully tetrahedral meshes results in fully tetrahedral refined meshes, while refinement of hexahedral meshes or combined hexahedral, pyramidal, prismatic and tetrahedral meshes results in refined meshes contain combinations of these types of elements. This work was presented at the AIAA Aerospace Sciences Meeting in Reno NV, in January 1997. The efficiency of the technique was demonstrated by generating a refined tetrahedral mesh of 1.3 million points on an aircraft configuration, and a refined hexahedral mesh of over 3 million points on an ONEAR M6 wing configuration using a SUN ULTRA workstation.

The adaptive meshing algorithm has been coupled to the flow solver. Additional improvements to the flow solver are required to enable consistently rapid multigrid convergence on highly adapted meshes.

Unstructured multigrid convergence acceleration for highly stretched meshes

For high-Reynolds number viscous flow simulations, efficient resolution of the thin boundary layer and wake regions requires mesh spacings several orders of magnitude smaller in the normal direction than in the streamwise direction. This extreme grid stretching results in poor multigrid convergence rates, usually one to two orders of magnitude slower than those observed for equivalent inviscid flow problems without grid stretching. The purpose of this work is to devise improved multigrid techniques for acceleration convergence on anisotropic grids of this type.

The approach taken consists of using implicit line solvers in the direction normal to the grid stretching combined with semi-coarsening or directional coarsening multigrid to alleviate the stiffness due to grid anisotropy. A graph algorithm has been implemented which combines edges of the original mesh into lines which follow the direction of maximum coupling in the unstructured mesh. A similar algorithm is used to selectively coarsen the original fine mesh by removing points along the directions of strong coupling, thus recursively generating a hierarchical set of coarse meshes for the multigrid algorithm. Preconditioning to alleviate low-Mach number induced stiffness in the equations has also been implemented. Finally, a Newton GMRES technique has been implemented which uses the above described multigrid and preconditioning techniques as a preconditioner themselves. Using this approach, convergence rates for viscous flows which are independent of the degree of mesh stretching can be obtained. These multigrid convergence rates are close to those achieved on equivalent inviscid flow problems with no mesh stretching.

At present, this methodology has been demonstrated using the "overset-mesh" multigrid technique, which relies on a sequence of non-nested unstructured grids. The methodology is to be implemented within the more practical "agglomeration" or algebraic multigrid strategy, firstly in two dimensions, and later in three dimensions.

R. A. NICOLAIDES

The use of divergence boundary conditions in computational electromagnetics

This project concerns the conditions under which divergence boundary conditions may be used in place of interior divergence constraints which appear in Maxwell's equations and the incompressible fluid flow. Our emphasis is on the electromagnetic case. There are a number of advantages in using the boundary condition form of the constraint. The most important one is that it becomes possible to use standard finite element subspaces instead of the far more complicated edge elements which are usually applied to computational electromagnetics. In addition, the divergence boundary condition is a natural boundary condition and that is another reason why the implementation is relatively straightforward.

We have obtained a fairly complete solution to the problem of when it is safe to use divergence boundary conditions. Briefly, it is necessary and sufficient that the domain is such that the Poisson equation has a solution in the Sobolev space H^2 . If this is satisfied, the use of the divergence boundary condition gives the same solution as if the full interior constraint were applied. Unfortunately, this implies that if a domain has reentrant parts, then the conditions are not equivalent and using the boundary formulation can (and normally will) cause approximation schemes to converge to an incorrect result. To repair this difficulty it is necessary to incorporate singular functions in the approximation, but it is not known what these functions are in three dimensions.

Future plans call for discovery of techniques which permit the use of the boundary conditions in more general situations. A number of modifications to the trial functions near singular points may be envisaged and are under investigation.

RONALD H. NOWACZYK

Perceptions on the effectiveness of engineering teams

While considerable work has been devoted to our understanding of effectiveness in management and business teams, little empirical work has addressed the functioning of science and engineering teams. Given that scientists and engineers are more likely today to work together as part of interdisciplinary teams, this research was undertaken in order to identify the critical factors and dimensions to successful teamwork on engineering and science design teams. The objective of this research is to identify not only those team factors that are common to both business and engineering teams, but also those factors unique to the mission and success of engineering and science teams.

The approach to this problem is twofold. One is the observation of engineering teams at work. The other involves interviews and surveys of scientists and engineers on their perceptions of effective teamwork. For the purposes of this study, a team has been defined as "a group of individuals working together toward a common goal, product, or solution that requires the sharing of expertise, knowledge, and ideas in a cooperative and interdependent fashion." To limit the focus to engineering and science teams specifically, the following definition of an engineering team has been used, "teams with an engineering design, process, or product focus." A paper-and-pencil survey has been developed and distributed to engineers and scientists at NASA-LaRC. The survey consists of three sections. The first two sections focus on behaviors observed in teams. The first section examines organizational culture and external influences on the team. The second section examines

behaviors internal to a team (e.g., communication, goal setting, conflict resolution). The third section contains demographic questions and includes items that provide a picture of the respondent as a team member. The survey takes approximately 30 minutes to complete. Preliminary analysis of the results of the survey are ongoing. This research has identified several factors related to the success of an engineering design team.

Ongoing work is directed at the identification of the essential features found in successful engineering teams. Additional work will focus on the development and validation of assessment instruments to be used by teams as part of a self-evaluation process of team functioning. These paper-and-pencil instruments can be used by teams during the life of the team and at the conclusion of the team's project. Later work will be devoted to the development of brief team exercises that can improve engineering team performance. These exercises will be designed so that they can be accomplished in less than one hour and can be administered by the team leader or an outside facilitator. Exercise objectives, guidelines for the administration of the exercise, and interpretation of the results will be included with each exercise.

CHI-WANG SHU

Discontinuous Galerkin method applied to viscous problems

Our objective is to study and apply high-order discontinuous Galerkin finite element methods for convection dominated viscous problems. The applications will be problems in aeroacoustics and other time-dependent flow problems with complicated solution structure.

Jointly with Harold Atkins at NASA Langley, we are investigating a framework of discontinuous Galerkin method applied to a first-order system by introducing auxiliary variables representing the first derivatives of the solution. Emphasis is put upon finding good local implicit procedure which can effectively enhance CFL stability limits. Jointly with Bernardo Cockburn of University of Minnesota, we have proved nonlinear stability and obtained error estimates for the linear case for this framework of the discontinuous Galerkin method.

Research will be continued for this class of high order discontinuous Galerkin methods and the applications. Related research on high-order finite difference and spectral methods will also be performed.

DAVID SIDILKOVER

Essentially optimal multigrid solvers for the flow equations

The main objective of this work is to develop a simple essentially optimal multigrid solver for the steady Euler equations (both incompressible and compressible), i.e., a solver whose efficiency is similar to one of a solver for Poisson or Full-Potential equations.

A simple discretization of the Euler equations (both compressible and incompressible) that facilitates such a solver was constructed recently. The discrete equations to be solved at each node are "assembled" from the residuals of the Euler system on the grid elements having this node as a common vertex. The elliptic factor separates in the form of Poisson (incompressible case) or Full-Potential (compressible case) operators acting on the pressure. The basic algorithm for incompressible Euler equations has been implemented in the framework of two engineering codes

(structured and unstructured grids) capable of treating complex geometries. The essentially optimal multigrid efficiency was already demonstrated in some complex geometry cases.

The current work includes addressing some more realistic test cases. Another issue that is being treated currently is incorporating genuinely multidimensional high-resolution upwind schemes to discretize the momentum equations in the context of unstructured triangular grids. The plans for the near future include generalizing the implementation to incompressible laminar Navier-Stokes and compressible Euler equations.

This is a joint work with Drs. T.W.Roberts and R.C.Swanson of NASA Langley Research Center.

Multidimensional upwinding in the control volumes context

A genuinely multidimensional high-resolution discretization for the steady compressible Euler equations was constructed recently. One of the fundamental advantage of this approach is that the genuinely multidimensional high-resolution mechanism (unlike the standard one) does not damage the stability properties of the scheme. The scheme was formulated in the residual distribution context. The existing codes, however, rely mostly on control volume type discretization. It will be useful, therefore, to formulate a genuinely multidimensional control volume scheme, since it can be easily incorporated into the existing codes.

A control volume type genuinely multidimensional scheme for the Euler equations was constructed. A standard first-order upwind scheme is used as a basic building block. Multidimensional high-resolution corrections based on a multidimensional limiter are added to the first order numerical fluxes. The resulting high-resolution scheme relies on a compact "9-point box" stencil. The preliminary numerical experiments verify the high-resolution properties of the scheme.

Since the multidimensional high-resolution corrections (unlike the standard dimension-by-dimension ones) do not damage the stability properties of the scheme, the immediate practical benefit may be due to the possibility of using a simple Gauss-Seidel relaxation as a smoother in the multigrid solver. Also, a recent study revealed that some variants of genuinely multidimensional control volume scheme are factorizable. The latter property is crucial for the study concerned with achieving the optimal efficiency of a multigrid solver.

RALPH C. SMITH

PDE-based control of thin shell dynamics

The problem of controlling structural dynamics modeled by thin shell equations arises in applications ranging from noise control in a fuselage to flow control in flexible pipes. A mechanism common to all such applications is the coupling between component displacements due to curvature and geometry. The use of models which quantify the coupling are crucial to the success of the controller. The objective of this investigation was the development and numerical implementation of PDE-based linear quadratic regulator (LQR) control techniques for applications involving thin cylindrical shells.

Donnell-Mushtari thin shell equations modified to incorporate passive and active piezoceramic patch contributions were used to model the shell dynamics. This model characterizes the coupling and physics in a large number of thin shell applications and provides a setting from which analysis,

numerical approximation techniques and control methods can readily be extended to more complex geometries or coupled systems. The model dynamics were approximated using a Galerkin formulation with a spline basis employed in the axial direction and Fourier expansions used in the periodic circumferential direction. A convergence framework for the LQR control problem was developed and numerical control algorithms were implemented. Initial numerical studies demonstrated that 80-90% reductions in displacement levels can be obtained when voltages computed in this manner are fed back to controlling patches on the shell.

In future investigations, we will consider the problem of combining the full state LQR method with a state estimator to obtain an output control method which requires only a limited number of state observations. The development of such state estimators and output control methods is necessary before the methods can be experimentally implemented.

This research was conducted in collaboration with R.C.H. del Rosario (North Carolina State University).

V. VENKATAKRISHNAN

Design optimization

Work in design optimization has been completed in two dimensions using unstructured grids. A design capability for two-dimensional Euler and laminar Navier-Stokes equations has been developed.

In the earlier approach, the adjoint for the viscous terms were derived using a discrete approach, whereas the discretization of the continuous adjoint equation for the Euler equations were worked out in such a way that there is complete correspondence with the discrete adjoint only for a first order accurate scheme. A higher order discretization was achieved by reconstruction. In order to further improve the accuracy of the derivatives (compared to the finite difference derivatives) a discretization of the adjoint equation has been implemented that is exactly equivalent to the discrete adjoint method for first and second order discretizations, both for incompressible and compressible flows. It is felt that this step was needed in addition to grid sensitivity terms in order to obtain proper derivatives for gross relative motions, such as translation and rotation of airfoil elements relative to one another. One of the important findings in this paper was that not all cost functions are admissible. The terms arising from taking variations of cost function have to be balanced by terms arising from the residual operator. Thus, for inviscid flows, only cost functions involving pressure are admissible. For Navier-Stokes, the cost functions have to involve the entire viscous tensor. Thrust lift, drag and moments are admissible, provided the viscous contributions are also taken into account. For viscous flows, pressure matching would appear to be inadmissible, but we were able to show that it is indeed a proper cost function.

The findings were presented as a technical paper at the Aerospace Sciences Meeting in Reno, January 1997. Extensions to compressible turbulent viscous flows and three-dimensional applications are planned.

This work is being done in collaboration with W. K. Anderson of NASA Langley Research Center.

BRAM VAN LEER

Local preconditioning for the Navier-Stokes equations

Local preconditioning has proved a valuable tool for convergence acceleration of Euler and Navier-Stokes codes, when low-speed and/or transonic regions appear in the flow. Proper preconditioning also prevents loss of accuracy in low-speed regions. High-lift, V/STOL and propulsion flows are examples where compressibility cannot be ignored, so that compressible-flow codes must be used; the embedded low-speed regions then slow down convergence. High-lift flows are of interest to NASA LaRC under the Advanced Subsonic Technology program.

Local preconditioners for the Navier-Stokes equations, based on a dispersion analysis, look promising as a means to speed up convergence of both single- and multi-grid processes, and to preserve the accuracy of the computed flow. However, they still suffer from loss of robustness near stagnation points, for low cell-Reynolds numbers, and for large aspect ratios. In contrast, Jacobi-type preconditioning is a good basis for multi-grid relaxation, but it has no single-grid benefit, and it does not prevent the accuracy loss in low-speed regions. Together with D. Lee (U. of Michigan) and W. Kleb (LaRC, Aerothermodynamics Branch) I have studied combinations of dispersion-analysis based and Jacobi-type preconditioning, under the assumption that a blend or embedding of these techniques may have all the robustness of the latter, without giving up the accuracy and single-grid advantages of the former.

It is still too early to tell whether this approach will lead anywhere. In the coming summer this research will be continued, in cooperation with Eli Turkel (Tel-Aviv U. and ICASE).

FLUID MECHANICS

P. BALAKUMAR

Non-linear equilibrium solutions and secondary instabilities

The transition process from a laminar to turbulent state in a shear flow involves essentially four steps: (1) Primary (linear) instability (2) nonlinear saturation of the primary instability and formation of a secondary periodic flow, (3) linear instability of this secondary flow (secondary instability) and (4) highly nonlinear or breakdown region. There exist three methods to investigate the nonlinear evolution of disturbances in a shear flow. One is the weakly nonlinear theory, the second is to solve the full Navier-Stokes equations as an initial value problem, and the third is to directly seek nonlinear equilibrium solutions as exact solutions of the Navier-Stokes equations. The objective of this work is to compute these nonlinear neutral solutions and to determine the nonlinear critical Reynolds numbers and their instabilities.

Two-dimensional nonlinear equilibrium solutions are computed for plane Poiseuille-Couette flows and for attachment line boundary layers. The equations are solved using the two-point fourth order compact scheme and the Newton-Raphson iteration technique. Starting with the neutral solution for plane Poiseuille flow, the nonlinear neutral curves for plane Poiseuille-Couette flow are mapped out by increasing the Couette velocity component. It is concluded that two-dimensional nonlinear equilibrium solutions do not exist beyond a critical Couette velocity. Similar computations for the attachment line boundary layer show that this flow is stable for all types of two-dimensional disturbances below the critical Reynolds number of 510.

In the future, we plan to investigate three-dimensional bifurcations from these equilibrium solutions. This requires efficient eigenvalue solvers for large sparse matrices.

ALVIN BAYLISS

Numerical simulation of jet flow/acoustics/structural interaction

The objective of this research is to simulate the coupling between sound from a jet with an array of flexible aircraft panels. In a real aircraft panel, vibration can lead both to increased interior noise levels and to structural fatigue. Our research is designed to provide a tool to simulate panel response due to realistic sound sources in the jet. This will provide a means of assessing the degree of panel vibration and provide a tool to assess different mechanisms to control panel vibration.

We have developed a numerical scheme to compute panel response and radiation in fully self-consistent manner (i.e., panel response and radiation are fully coupled to the jet acoustics and flow field). In previous work, we have simulated panel response for two-dimensional Cartesian jets under a variety of flow conditions, including low Mach number jets, jets in forward motion, high subsonic jets and supersonic jets. In each case, we have computed jet response, panel response and radiation. In our latest work, we have computed the response of the panels to sound from an ideally expanded supersonic jet. We have computed Mach wave radiation, demonstrated the effect of this radiation on panels, determined the effect of panel location on panel loading and vibration

and determined characteristics of the interior acoustic distribution. This work was presented at the AIAA Aerospace Sciences conference and an ICASE report on this work has been prepared.

We are in the process of extending our numerical model from two-dimensional Cartesian to nonaxisymmetric cylindrical coordinates. This will enable us to simulate realistic, three dimensional instability waves in a jet together with the ensuing sound generation and panel loading. Upon validation of this code, we plan to use it in a coupled manner with a code to predict the panel response to the jet sound loading.

This work is a collaboration between A. Bayliss and L. Maestrello of NASA Langley Research Center.

AYODEJI DEMUREN

Numerical simulation of complex turbulent jets

Turbulent jets are encountered in many aerodynamic flows and industrial applications. In many of these applications, rapid mixing is desirable, either to reduce noise from jet engines, to promote fuel-air mixing in combustion chambers, or to promote rapid dilution of pollutants exhausting into the atmosphere. Jets with non-circular cross-sections tend to experience more rapid spreading and mixing than circular ones. Although it is believed that the complex vorticity field is responsible for this, the mechanisms are not fully understood. The objective of the current research is to increase the understanding of jet mixing processes through direct simulation and large eddy simulation of jets with elliptical and rectangular cross-sections.

The computational scheme uses a low-storage third-order Runge-Kutta formulation for temporal integration and a fourth-order compact scheme for spatial integration of the incompressible Navier-Stokes equations for turbulent jets with rectangular cross-sections. To satisfy the divergence free condition for the velocity field, a Poisson equation is solved for pressure which must also be approximated by the same compact scheme for consistency. Simulations of rectangular jets at a low Reynolds number of 750 have been performed. The vortex dynamics shows some interesting mechanisms. Larger eddy simulations of rectangular, elliptic and circular jets at a nominally high Reynolds number of 75,000 have been performed. The simulation database is now available for post-processing. Some statistical moments have been computed, and some visualizations of the vortex dynamics have been made. The significance of the results is that we now know some processes for the generation of streamwise vorticity in jets with complex cross-sections. We also have a database which can be used for turbulence modeling and acoustic computations.

Further analysis of the results are to be performed, and the methodology will be developed for using the results as sources for acoustic computations using linearized Euler equations.

SHARATH S. GIRIMAJI

Non-equilibrium algebraic Reynolds stress modeling

Work on algebraic modeling of Reynolds stress away from the equilibrium state of turbulence is continuing. This model must be considered an algebraic approximation of the modeled Reynolds stress differential equation.

In strain-rate dominated flows, the solution of the Reynolds stress differential equation typically exhibits three distinct stages. In the *early transient stage*, the solution undergoes rapid change and

is a strong function of the initial condition. Due to the non-linearity of the evolution equation, the initial conditions are soon forgotten and the system reaches a state of quasi-equilibrium. In this *slow manifold* stage, the solution is a function of the dynamics of the equation only. Finally, the solution slowly evolves along the slow-manifold to the time-invariant equilibrium solution. It now appears that the slow manifold regime (and, of course, the equilibrium state) are amenable to algebraic description. This description is currently being sought.

In rotation dominated flows, the solution of the Reynolds stress differential equation behaves somewhat differently, going to a limit cycle, rather than to an attracting fixed point. It appears that the late time solution still might be a strong function of the initial conditions, precluding algebraic description. Further investigation is underway using techniques of non-linear dynamical system analysis.

Algebraic modeling of scalar flux

Modeling of turbulent scalar (and temperature) flux is very important in computation of turbulent reacting flows. The objective of this project is to investigate various terms in its evolution equation using DNS data and to evaluate the potential for deriving an algebraic model.

Direct numerical simulation data of Rayleigh-Bernard convection was used to evaluate the modeling assumptions pertaining to algebraic scalar flux modeling. Also standard pressure-strain and pressure-temperature gradient correlation models for buoyant flows were evaluated using numerical data. It appears that the standard models as well as the algebraic assumptions are inadequate, especially near the edges of the flow. It is possible that the agreement would be better in higher Reynolds number flows for which the models are developed. New pressure correlation models that improve agreement with data have been developed in this study. The ensuing algebraic model is still not very accurate, but reasonably adequate away from the boundaries of the flow.

Further work is planned to extend this analysis to mixed convection flows and to continue development of improved pressure correlation models.

This research was conducted in collaboration with S. Balachander, Univ. of Illinois, Urbana Champagne.

Equilibrium states of homogeneous turbulence

The objective of this study is to better understand the non-linear dynamical system of equations that govern the Reynolds stress evolution in homogeneous turbulence.

The study is in its incipient stages. First we intend to perform a fixed point analysis of the modeled Reynolds stress evolution equations using representation theory. Then we will attempt to characterize the non-equilibrium behavior dictated by the evolution equations.

This study is expected to impact three areas of turbulence model development: (i) towards developing new pressure-strain correlation and dissipation models that are more consistent with Navier-Stokes physics than present models; (ii) towards developing improved algebraic Reynolds stress models, especially away from equilibrium; and (iii) to better understand the behavior of current models, so that designers (users) can apply the models more effectively and intelligently.

C.E. GROSCH

Simulation of mixing enhancement in hot supersonic jets

Two programs have been initiated by NASA that have identified engine noise reduction as an enabling technology. These programs are the NASA High Speed Research (HSR) and Advanced Subsonics Technology (AST) programs. In the HSR program, jet noise is the principle contributor. In the AST program, jet noise is the principle contributor for aircraft in the current fleet, where a goal of 3dB reduction is established for aircraft engines with bypass ratios up to 5. Methods used to reduce jet noise in both programs often utilize concepts that enhance mixing between high and low speed streams. Of these, the most popular methods utilize concepts that introduce streamwise axial vorticity. Experimental observations show that the presence of small tabs on the edge of a hot, compressible jet exiting into a slower moving, colder ambient flow can increase the rate of spreading of the jet. This suggests that the axial vorticity introduced by the tabs increases the rate of mixing of the jet and the ambient fluid. The objective of this research is to show that numerical simulations can model these effects and to elucidate the physical mechanism responsible for the increased mixing and spreading rate of such hot supersonic jets.

A set of calculations was carried out using the compressible three-dimensional Navier-Stokes equations for a rectangular jet in a square channel. First the flow was simulated without the tabs, obtaining reasonable agreement with experimental measurements of the velocity. Then the flow was simulated without tabs over a range of values of the convective Mach number in order to determine the dependence of the mixing on this parameter. Simulations with modeled tabs were also carried out. In these calculations the effect of the tabs on the flow was modeled by pairs of counter rotating vortices. Calculations were carried out using from one to eight tabs. Both "necklace" and "trailing" type vortices can be simulated, depending on the sense of rotation of the model vortices. Calculations of the flow with these present showed that the tabs increased the thickness of the jet by about 25% compared to the flow without the tabs. A mixing parameter was increased by a factor of about 2.5 by using six tabs. The results of the calculations also elucidated the basic physical mechanism of the interaction of the vortices generated by tabs with a hot jet as well as the mechanism causing the increased jet thickening and increased mixing. In short, the streamwise vortices transport the hot, higher momentum fluid from the central region of the jet to the colder, lower momentum region of the coflow and vice versa. This increases the z -ward transport of x momentum as well as increasing the mixing of hot and cool fluid. The results also show that the initial configuration of "necklace" vortices on both the upper and lower edges of the jet is generally unstable while that of "trailing" vortices is stable. Finally, there was good qualitative comparisons with flow visualizations. As a result of this study, the physical mechanisms are understood and the use of simulation was validated as a tool for further exploration of configurations and for the effects of varying flow parameters.

Future work will include simulation of jets with non-rectangular cross-section and lobe injectors, and parametric studies of the number and placement of the tabs.

This research was carried out in collaboration with J.M. Seiner (Aeroacoustics Branch, NASA Langley Research Center), M.Y. Hussaini (Florida State University) and T.L. Jackson (ICASE).

M. EHTESHAM HAYDER

Parallelization of an aeroacoustics code

The objective of this effort is to compute noise fields using an aeroacoustics code on a parallel computer.

We compute the noise field using the discontinuous Galerkin method. High numerical accuracy along with the use of unstructured grids make this code suitable for aeroacoustic computations for problems with complex geometries. This code is written in C++ and we plan to use message passing library calls to parallelize it.

We plan to examine various issues related to parallelizing this application and its extensions in future. This work is being done in collaboration with Harold Atkins (NASA Langley).

Non-reflecting boundary conditions

The treatment of boundary conditions is important for the accurate numerical simulations. Our objective is to examine and formulate boundary conditions to minimize numerical reflections.

We examined the effectiveness of buffer layers as non-reflecting computational boundaries for the Euler equations. The absorbing-layer equations were constructed by operator-splitting of the governing equations in the coordinate directions and introducing absorption coefficients in each split equation. This methodology follows that of the Perfectly Matched Layer introduced by Berenger for the numerical solution of Maxwell's equations. We considered the application of absorbing layers in a few physical problems, such as shock-vortex interactions, a plane free shear flow and an axisymmetric jet with emphasis on acoustic wave propagation. We also examined the effectiveness of the absorbing layer for the solution of the Euler equations where nonlinear effects are significant.

This investigation is being done in collaboration with Fang Hu (Old Dominion University/ICASE) and M.Y. Hussaini (Florida State University).

TOM JACKSON

Algebraic instabilities in Blasius boundary layer

Algebraic instabilities arise when some initial disturbances, owing their presence to a finite level of noise present in any flow, grow sufficiently to trigger nonlinear mechanisms or to provide new basic states for secondary instabilities. These instabilities are distinguished from exponential instabilities, where infinitesimal disturbances always grow exponentially in time. The presence of algebraic instabilities may lead to the so-called "bypass mechanisms". Work is continuing on the evolution of disturbances in the Blasius boundary layer flow. This work offers a means whereby completely arbitrary initial input can be specified and the resulting temporal behavior, including both early time transients and the long time asymptotics, can be determined.

The approach taken to explore the transient dynamics include linear theory and DNS calculations. The bases for the linear analysis are: (a) linearization of the governing equations; (b) Fourier decomposition in the spanwise and streamwise directions of the flow and; (c) numerical integration of the resulting partial differential equations. DNS spatial calculations are performed to compare with linear theory and to explore the nonlinear dynamics. Because the DNS calculations are expensive when trying to compare to linear theory, we are currently writing a 2D version using high

order compact schemes to bridge the gap between the fully nonlinear DNS calculations and the calculations of the linear theory. The results provide explicitly both the early time transients and the long time asymptotic behavior of any perturbation. With this knowledge it is then possible to devise means for flow control and it is possible to either delay or enhance disturbances as the need may be. In addition, the important problem of receptivity can also be analyzed within this framework. All linear results are compared to the equivalent spatial problem using DNS.

In the future, we plan to investigate the receptivity of the boundary layer, as well as various possible control mechanisms for the boundary layer.

This work is conducted in collaboration with Ronald Joslin (NASA), William Criminale (University of Washington) and D. Glenn Lasseigne (Old Dominion University).

LI-SHI LUO

Theory of Lattice Boltzmann equations

Recently, the method of Lattice Boltzmann equation (LBE) has attracted the attention of physicists, chemists, engineers, and computational scientists working in broad spectrum of scientific disciplines, in part because the LBE method possesses the following features:

- The LBE algorithm is intrinsically parallel;
- The LBE algorithm is rather simple, thus the programming effort is minimal;
- The LBE provides a consistent framework of thermodynamics at the kinetic level. Interactions (potentials) can be easily modeled.
- The LBE method has broad applications in hydrodynamics, magneto-hydrodynamics, and complex fluids.

Despite the rapid development, the theory of the LBE method has not yet been formulated in a rigorous fashion.

Recently, we were able to give an *a priori* derivation of LBE from its continuous counterpart. This work has placed LBE on a firm mathematical foundation.

We intend to extend the work to LBE models of thermo-hydrodynamics, multi-phase and multi-component fluids in the future.

This research was conducted in collaboration with Dr. Hudong Chen (Exa Corporation), Prof. George Vahala (College of William & Mary), Dr. Xiaoyi He (Los Alamos National Laboratory), and Dr. Nick Martys (NIST).

Direct numerical simulations by LBE method

Because the LBE method is relative new, there are not many numerical benchmarks being done. There is an urgent need to conduct high quality direct numerical simulations by LBE methods.

By using an incompressible LBE model of He and Luo, we have done simulations of the two-dimensional lid-driven cavity flow, which has been a paradigm to test CFD codes. Our results quantitatively agree with previous ones by finite difference, and finite element methods.

We propose to use the LBE method to simulate other flow systems, such as sudden expansions and the back-facing step.

This research was conducted in collaboration with Exa Corporation, Ford Motor Company, and Prof. George Vahala (College of William & Mary).

JAMES E. MARTIN

Swirling jets under helical perturbation

Swirling jets are of importance in a variety of technical applications involving mixing, propulsion and combustion processes. Under certain conditions, swirling jets are also known to produce vortex breakdown events. Although much is known about the coherent structures and the dynamics of jets without swirl, very little is known about jets with an additional swirl velocity component. In an earlier study, we investigated swirling jets subject to axisymmetric and azimuthal perturbations. The results of that study revealed a complex nonlinear interaction of several competing instability mechanisms when an additional swirl velocity component is included. It was determined that with swirl, a series of counter-rotating vortex rings may replace the standard corotating vortex ring structure of jets without swirl. The aim of the present investigation is to determine the large scale structures and their dynamics for swirling jets under helical perturbation.

We utilize Lagrangian vortex filament simulations to study the jet's nonlinear evolution. Using helical filaments, the cylindrical shear layer is given both circumferential and streamwise vorticity. By changing the pitch of the helices, different jumps in the axial and circumferential velocity across the shear layer may be considered. In the case of axisymmetric perturbations, we were able to model the swirl velocity within the jet using a nondeformable axial line vortex. To consider helical perturbations, we must replace the axial line vortex with a deformable vortex filament. This modification to the code has been completed and we are now ready to consider the helical perturbation case.

As a preliminary calculation we will consider an $m = \pm 2$ helical perturbation. An additional azimuthal perturbation will then be included to break the helical symmetry of this calculation. This research was conducted in collaboration with E. Meiburg (University of Southern California).

J.R. RISTORCELLI

Aeroacoustics

A low fluctuating Mach number analysis has been used to obtain the variance of the dilatational from turbulence closures. A similar analysis is being used to obtain the variance of the dilatational *rate* - which has been argued by Ribner (1962) to be the acoustic source. Low Mach number asymptotics have been used to formally vindicate Ribner's contention. The procedure will allow prediction of the near field broadbanded sound source intensity from RANS type turbulence models. The work at this point is directed towards a parameterization of the broadbanded sound source intensity field in terms of the turbulence statistics. Issues dealing with the far-field acoustic intensity cannot be addressed at this stage of the work: the required statistics for the propagation problem are not available from single-point closures. Profiles of the variance of the dilatation have been calculated in flows over aerofoils and for the shedding cylinder by M. Sanetrik (Analytical Services and Materials, Inc.)

The main impediments to the rational application of the low Mach number asymptotics relating the broadband sound-source statistics to the turbulence statistics is a useful and kinematically

correct representation for the two-point correlation function. This is a present area of research that needs to be addressed before one can rationally apply the expressions for the variances of the dilatational field.

The work is being carried out in conjunction with one of the aeroacoustical projects of Aerodynamic and Acoustics Methods Branch. The results of the analysis are being implemented, by E. Hayder (ICASE) in free shear (jet) flow for which there exists an abundance of data.

Compressible turbulence modeling

Primary concern is with the creation of a consistent set of models for the effects of compressibility. A systematic perturbation procedure applicable to low turbulent Mach number flows in "compact" flow domains is used. The analysis applies to turbulence in supersonic and hypersonic flows which have a low turbulent Mach number.

Focus has been on the reversible transfer due to the pressure-dilatation. The theoretical aspects of the effects of the pressure-dilatation are complete. A very good comparison of the scalings predicted by the analysis and those seen in the DNS have been found. The theory appears to have been unquestionably vindicated. The results of the analysis have been used to compute the mixing layer and decreases in mixing layer growth are observed. Additional and more substantial decreases in the mixing layer growth are expected to come from the reductions in the shear stress anisotropy which appears to be a more important suppression mechanism in simple compressible shear flows. This effect is associated with the pressure-strain correlation. Theoretical progress has been made at the second-order closure level to obtain the effects of compressibility on the shear anisotropy. Further results will await the completion of a lengthy mathematical development.

Future plans in this area include the completion of the analytical work for the pressure-strain. The work is at a second-order closure level to account for reduction of the shear anisotropy by compressibility. Work continues on understanding the pressure-dilatation results using the recent DNS by Blaisdell which is formulated with consistent initial conditions. Also intended is an investigation of the compressible dissipation scalings that result from the analysis. In addition to further evaluation of the weakly compressible theory this has implications for aeroacoustical problems that presently form the major portion of our effort.

Some of the theoretical developments in compressible turbulence are being made possible by collaboration with G. Blaisdell (Purdue University).

ROBERT RUBINSTEIN

Rotating turbulent shear flow

The goal of this work is to develop a two-equation model for rotating turbulent shear flows. Such flows occur in turbomachinery and in geophysical applications. Two-equation modeling of rotating turbulence is especially challenging because rotation does not explicitly enter any of the exact transport equations. It provides a simple but important special case of the effects of external agencies on turbulent flows.

The modeling is based on the two-point statistics of rotating turbulence. It is shown that some previously developed heuristic descriptions of strongly rotating turbulence can be justified as asymptotic solutions of the equations of the direct interaction approximation. The two-point

descriptors provided by this analysis can be substituted in expressions for the eddy viscosity, developed by Leslie, Yoshizawa and others, and in expressions for the production and destruction of dissipation rate, developed in the course of this investigation. The two-equation model which results clearly shows the suppression of energy transfer which is characteristic of rotating turbulence. This approach has also been applied to derive a modified Smagorinsky sub-grid scale viscosity for use in large eddy simulations.

Efforts to test and refine this model by comparing with experimental data on rotating pipe flow are continuing.

This work is done in collaboration with Ye Zhou.

Simulation methods for sound radiated by turbulence

The goal of this work is to model the sound radiated by sub-grid scale motions in large eddy simulations of turbulent flows. Large eddy simulations are an increasingly widespread approach to computational aeroacoustics and generally predict the peak in the acoustic spectrum of turbulence generated sound quite well. But this approach necessarily suppresses the contribution of the fine-grained sub-grid scale motions to the radiated sound.

The problem is approached by developing a stochastic simulation for the sub-grid scale motions based on commonly accepted models for the two-point two-time correlation function of turbulence. The parameters in this function are available from the resolved scales of motion.

Previous large eddy simulation of turbulent jets with non-circular cross section will be used to demonstrate this method. The resolved fluctuations will be enhanced by sub-grid scale fluctuations synthesized as described above. The resulting fluctuating field will be used as a sound source to predict the far-field radiated sound.

This work is done in collaboration with Ye Zhou and Ayodeji Demuren.

Plasma drag reduction

The goal of this work is to test coupling between acoustic and ion-acoustic waves as a mechanism for suppression of shock waves in weakly ionized gases, a phenomenon proposed on the basis of various experimental observations.

A simple model for waves in a medium with two sound speeds has been developed in which the wave modes are coupled through collision-like terms. Analysis of the dispersion relation in this medium show that long acoustic waves, which correspond to the small wave speed, propagate faster as a result of coupling whereas long ion-acoustic waves, which correspond to the high wave speed, are over-damped. But to achieve the increased sound propagation speed proposed to explain the experimental observations, much larger coupling constants than expected are needed. These conclusions are supported by numerical simulations of the propagation of pulses through this medium.

The effect of this type of coupling on finite amplitude waves will be analyzed by numerical simulation of one-dimensional shock waves in this medium. The anomalously high coupling will also be studied further: a medium with two sound speeds must be in a steady state far from thermal equilibrium. Like turbulence, this steady state is maintained by a balance between a source and sink of energy. It is possible that higher than expected coupling, the result of ion densities larger

than the input ion density, are required to maintain this steady state and that the energy source which maintains the increased ion density is the mean flow.

This work is done in collaboration with Aaron Auslender (NASA Langley).

SIVA THANGAM

Analysis of turbulence models for complex flows with curvature and rotation

Fluid flows encountered in most engineering applications are quite complex and often include significant levels of curvature and/or rotation. Turbulence models for such flows are predicated by the need for accurate estimates of turbulence stresses. The focus of research during the current phase involves the development of efficient turbulence models that can capture the effect of curvature and rotation as well as their validation. The work performed during the current semi-annual period involves collaborative efforts with Dr. Ye Zhou of ICASE and Professor K.R. Rajagopal of Texas A & M University as well as others at Stevens Institute of Technology and ICASE.

A turbulence model for rotating flows has been developed based on the phenomenological treatment of rotation modified energy spectrum. While the solid body rotation influences the energy transfer process significantly, it does not enter the kinetic energy equation explicitly. A dissipation rate equation with rotation rate dependent model coefficients has been developed. The model has been calibrated using several test cases and is currently being applied to rotating duct flows using a recursion renormalization group (r-RNG) theory based closure approach. In addition, the model for the effect of rotation and curvature will be also implemented with a generalized turbulence closure based on extended thermodynamics considerations.

Investigations of the development and implementation of two-equation models based on the algebraic representation of a recently developed second-order closure for compressible turbulence were also undertaken during the current period. The model utilizes both pressure-dilatation and dilatational dissipation for kinetic energy and has been successfully applied for wall- /bounded flows and supersonic mixing layers in conjunction with r-RNG theory based closure approach. This effort will be continued to include wake flows and other complex turbulent flows of aerodynamic relevance.

L. TING

Turbulent Boundary Layer, Acoustic and Structural Interaction

The interaction of panel oscillation with the boundary layer and incident acoustic waves simulating jet noise is essential for the prediction and control of panel fatigue and the transmission of external noise through panels of an airframe into the interior.

We study the implications of experimental data obtained by Dr. Lucio Maestrello (NASA Langley) on the interaction of incident acoustic waves, turbulent boundary layer and panel oscillations and the effectiveness of an active control device on the panel oscillation and the transmitted wave. It was found that the peak levels of the power spectra of panel response and transmitted wave increase by an order of 20dB in the presence of an incident pure-tone sound. When the level of the pure-tone sound exceeds a threshold value, the characteristics of the panel response change from periodic to chaotic. Simple mathematical models simulating certain aspects of the complex phenomena are being formulated. For numerical simulations, we examined a recent numerical program

developed by Dr. A. Frendi for the interaction of a flexible panel and a turbulent boundary layer in a supersonic stream. We found that the nonlinear panel equation missed a coupling term with the boundary layer wall shear. The missing term is being added to the program and the parameter characterizing its importance identified.

We intend to construct weakly nonlinear solutions of our simple model equations to explain how a single controller can effectively damp nonlinear oscillation of a panel near resonance and find the optimum amplitude, period and phase shift for the controller. We shall add the missing term in the numerical simulations of nonlinear panel oscillations coupled with turbulent boundary layer and acoustic excitation and show the importance of the wall shearing stress.

This research has been conducted in close collaboration with Dr. Lucio Maestrello (NASA Langley). To develop the numerical program for the simulation of the boundary layer/acoustic/structure interaction and control, active assistance of Dr. A. Frendi (Analytical Services and Materials, Inc.) is needed.

GEORGE VAHALA

Thermal Lattice Boltzmann (TLBE) simulations of variable Prandtl number turbulent flows

With the advent of multi-parallel processor machines, TLBE algorithms have become an interesting alternative to regular CFD codes for solving turbulence in complex flows. This is because TLBE codes are very simple (well over an order of magnitude shorter than CFD codes), easily extended to 3D and can readily handle complex flows. Moreover, since TLBE codes typically involve only free-streaming on a given lattice and relaxation at each lattice site, all these operations are purely local and so are ideal for multi-PE's. Moreover, since the form of the macroscopic conservation equations rely only on some symmetry properties of the collision operator, one typically chooses a kinetic equation that is as simple as possible.

Most LBE (and TLBE) approaches have used the BGK collision operator. However this is only a single-time relaxation operator and so the viscosity and thermal conductivity transport coefficients are intrinsically linked. This means that the Prandtl number is invariant for all flows. Using an extension to the single-time relaxation BGK operator proposed by Hudong Chen, we have been able to simulate variable Prandtl number flows. This new collision matrix has an extra parameter in the off-diagonal elements, it is readily invertible so the Chapman-Enskog analysis proceeds as simply as before. We have performed two-dimensional simulations Reynolds number = 35,000 on a 512 x 512 grid and shown the effect of variable Prandtl number on the coalescence of the vortices. Also, since TLBE methods are explicit, they are prone to numerical instabilities. We have performed a stability analysis on two TLBE codes: both 13-bit codes, one on a hexagonal lattice and the other on a square lattice. The maximal achievable Reynolds number was computed and it shows that hexagonal grids have much better stability properties than the square lattice.

Future work involves the optimization of the parallelization for the T3E, using MPI. With Li-Shi Luo (ICASE), we will begin studies of using more appropriate choices for the relaxation distribution function in an effort to cure the numerical instability of TLBE codes. Also there will be a consideration of the effects of spatial interpolation when non-uniform spatial grids are employed.

This work was done in collaboration with Li-Shi Luo (ICASE).

Turbulence modeling of the Toroidal wall heat load due to shear flows over cavities in the neutral gas blanket divertor regime

Projected heat loads onto the divertor plate in reactor tokamaks are a major cause of concern. One approach has relied on plasma recombination resulting in a cold (1 eV) neutral gas blanket between the divertor plate and the plasma flame front in the scrape-off layer (SOL). With plasma recombination and neutral gas ionization within the divertor region, there is strong coupling between the plasma dynamics and the neutral fluid flow. Since the plasma flows principally along the magnetic fields lines, the neutral flow will also be along the field lines, readily reaching Mach numbers > 1 and Reynolds numbers > 1000 . Since turbulence can be triggered in channel flows with $Re > 650$, it is of interest to see how neutral fluid turbulence will reduce the heat loads on the divertor plate.

We have restricted ourselves to considering the effect of 3D mean neutral velocity flow over toroidal cavities on a heat front as it propagates towards the divertor plate. The code ISAAC has been employed to study both the K-epsilon and the Algebraic Stress Models (ASM). It has been found that as the heat pulse propagates towards the divertor plate, the maximum heat load to the toroidal wall is a factor of 5 greater in ASM over K-Epsilon, which itself is a factor of 3 greater than that from laminar flow. More interesting is the result that in the wake of the heat front, the toroidal heat load for ASM is over a factor of 2 greater than the maximum K-Epsilon heat load. Both K-Epsilon and Laminar toroidal heat loads damp out rapidly in the wake of the heat pulse. These results are attributed to the effects of nonlinearities in mean vorticity and mean rate of strain for the off-diagonal Reynolds stress tensor components. These results are very encouraging as the resulting heat load on the divertor plate will be significantly reduced.

ISSAC, in its present form is not parallelized. Attempts underway at parallelization have unearthed f90 compiler bugs (work performed at NERSC) - but the hope is to speed up the code by a factor of 10. This is necessary since the three dimensional computations are quite massive with an executable over 200 MW. We will also be trying to pose the steady state problem, rather than the simpler initial value problem. This is a nontrivial extension.

This work was done in collaboration with Joseph Morrison (Analytical Services and Materials, Inc.)

YE ZHOU

On the universality of the Kolmogorov constant in numerical simulations of turbulence

Inertial-range behavior as postulated by Kolmogorov's similarity hypotheses is widely regarded as a fundamental characteristic of turbulence at high Reynolds number. In particular, in the inertial range of intermediate scales the one-dimensional longitudinal energy spectrum is given by

$$E_{11}(k_1) = C_1 \epsilon^{2/3} k_1^{-5/3} \quad (1)$$

where k_1 is the longitudinal wavenumber, C_1 is known as the Kolmogorov constant, and ϵ is the mean dissipation rate. The value of C_1 in this "five-thirds" scaling law has long been believed to be universal, with substantial experimental support. Recently, however, there is renewed debate on the universality of C_1 , in part because of new measurements at high Reynolds number (Praskovsky and Oncley) and a subsequent new similarity theory (Barenblatt and Goldfield) that appeared to suggest

a persistent Reynolds number dependence even at high Reynolds numbers. On the other hand, the conclusion from a new and very extensive survey of experimental data by Sreenivasan is that, taken collectively, measurements do not support such a dependence for the Kolmogorov constant. The value of C_1 averaged over many experiments cited by Sreenivasan is about 0.53 ± 0.055 , although some corrections for the estimation of dissipation using local isotropy assumptions in experiments may be warranted. This work is motivated by the survey of experimental data noted above, and will focus on similar issues arising in numerical simulations of isotropic turbulence.

We have presented new results on the Kolmogorov scaling of energy spectra and structure functions in the inertial range, from direct numerical simulations of stationary isotropic turbulence ranging from $R_\lambda = 38$ (on a 64^3 grid) to $R_\lambda = 240$ (on 512^3). It is pointed out that a plateau at $k\eta \approx 0.1 - 0.2$ commonly used to infer the Kolmogorov constant from the compensated three-dimensional energy spectrum does not represent proper inertial range behavior. Instead, a proper (if still approximate) inertial range emerges at $k\eta \approx 0.02 - 0.05$ when R_λ increases beyond 140. We find that the proportionality constants C and C_1 in the three- and one-dimensional compensated energy spectra are about 1.62 and 0.60 respectively. These values are closer to experimental data than reported in most previous numerical simulations. In particular, if isotropy relations are used then we may infer from the three-dimensional spectra that $C_1 = 18/55C \approx 0.53$, in excellent agreement with experimental data. However, the ratio C/C_1 in our results differs from the theoretical value of $55/18$, because of significant departures from isotropy in the computed spectra at the wavenumber range where inertial-range behavior is otherwise reasonably well approximated. Results on second- and third-order structure functions over a range of Reynolds numbers further suggest that the simulation database that we have accumulated should be useful for investigating other aspects of similarity scaling and Reynolds number dependence.

It is perhaps worth noting that the highest-resolution (512^3) simulation reported is at a higher Reynolds number and extends over a greater number of large-eddy turnover times than other studies reporting 512^3 results for stationary isotropic turbulence. The spectra and structure functions obtained from these and the present simulations demonstrate that, with the latest advances in massively parallel computing, issues in inertial-range similarity can now be addressed in a more reliable manner than possible before.

APPLIED COMPUTER SCIENCE

GIANFRANCO CIARDO

Automated parallelization of discrete state-space generation and analysis

A rich variety of modeling and analysis paradigms involve generating a discrete state-space, and then analyzing that space to determine performance and reliability measures. Unfortunately these state spaces can grow quite large, limiting the applicability of a discrete state-space analysis approach. Our objective is to increase the size of the models that can be so studied by using multiple processors to accelerate the generation and solution time, and the larger memory of the distributed system.

We have developed heuristics for automatic load balancing when generating a state space generation on a distributed network of N workstations or a N -processor system. The main idea is to partition the state space \mathcal{S} into $M \gg N$ classes, and then assign each class to one of the N processors, attempting to maintain an equal load. However, difficulties arise because we do not know beforehand the size of the M classes, so assigning them uniformly to the N processors is usually not a good strategy. The approach is to take a corrective action (the reassignment of one or more classes of states from an overloaded processor to an underloaded one) when the load is detected to be unbalanced at runtime. Two criteria can be used to decide that the load is unbalanced. The first criterion simply attempts to equalize the memory requirements over the N processors, based on a snapshot of the current memory usage. The second criterion instead attempts to equalize the future execution load over the N processors by examining the size of their queues of unexplored states. From our prototype implementation we have been able to gather evidence showing that our approach is quite resilient, without having to ask for user input in deciding how to partition the state space. Almost linear speedups have been achieved on an SP2.

In the future we plan to explore dynamic splitting of classes (as opposed to our current dynamic reallocation of statically defined classes), and to experiment with alternative multi-phase exploration strategies that will further allow us to increase the size of the state space being considered, with negligible performance overhead.

This research was conducted in collaboration with David Nicol (Dartmouth College).

Simulation of fluid stochastic Petri nets

When dealing with models having hundreds or thousands of objects flowing through them (messages, customers, parts, etc.), exact discrete-state approaches are infeasible, except under special structural conditions: the size of the state space is simply too large. An alternative is then to approximate a large population with a continuous quantity. This has led to the definition of Fluid Stochastic Petri Nets (FSPNs), where the state is defined as a tuple $(m_1, \dots, m_d, x_1, \dots, x_c)$, with m_i the number of tokens in discrete place p_i and x_j the fluid level in continuous place q_j . However, the numerical study of FSPN models is quite cumbersome, requiring the solution of systems of coupled differential equations. Previous proposals have been limited to models having only a handful of discrete states (m_1, \dots, m_d) and severe restrictions on the possible evolution

of the continuous components (x_1, \dots, x_c) . The objective of this work is to utilize discrete event simulation for the solution of these models.

One immediate advantage of our approach is that many restrictions on the stochastic behavior become unnecessary, hence the modeling power of the formalism is greatly enhanced. The continuous component of the state still presents difficulties beyond that of ordinary discrete-event simulation, but these can be managed with various levels of sophistication, depending on the type of stochastic behavior. In the most general case, we must solve a (small) system of ODE at each step of the simulation. If we can assume an uncoupled behavior for the flow rates in and out of each continuous place (if the overall flow rate for place q_j depends only on the discrete marking (m_1, \dots, m_d) and on the fluid level x_j), we must solve independent ODEs at each step of the simulation. An even better situation arises if the flow rates have piecewise-constant behavior (if the level in each fluid place increases or decreases at a constant rate until some threshold condition is reached, such as some fluid place becoming empty or full, or until a transition fires); in this case, no ODE solution is required, we simply need to schedule additional events corresponding to the hitting of a threshold. The ability to simulate this type of mixed discrete-and-continuous state models will open the door for extending the class of models that can be used to assess the performance of realistic systems, especially those with very large population of identical objects.

Future work will focus on efficient implementation of our findings, a highly non-trivial effort, due to the novelty of the approach.

This research was conducted in collaboration with David Nicol (Dartmouth College) and Kishor Trivedi (Duke University).

Kronecker operators for the description and solution of Markov models

The solution of the large Markov models arising in performability studies is limited by the size of the state space \mathcal{S} and of the transition rate matrix R . Even when storing R in sparse format, the practical limit for a modern workstation is substantially less than 10^6 states and 10^7 nonzero entries, if we want to avoid the use of virtual memory. Our objective is to increase by at least one order of magnitude these limits through the use of Kronecker, or tensor, operators for the implicit storage of R .

We decompose the model into K interacting submodels, each described by a (small) transition rate matrix $R^k, k = 1, \dots, K$, and express R as the Kronecker sum of the matrices R^k , plus some Kronecker products of “corrective factor” matrices, also of small size. Unlike previous approaches, we focus on methods that do not require vectors of size of the “potential state space,” which can be much larger than the already large actual state space. This requires a coordinated effort in discrete data structures/algorithms and in numerical methods, in order to reduce both the memory and the execution complexity of the solution process. We have been able to efficiently store R (using negligible memory) and \mathcal{S} (using little over $|\mathcal{S}|$ bytes, for arbitrarily large state spaces), while, at the same time, substantially reducing the computational overhead due to the use of Kronecker operators, from $\log |\mathcal{S}|$ to $\log |\mathcal{S}_K|$, the size of the local state space for one of the submodels. We have also investigated specialized methods, “perfect and imperfect partitioning,” that can be applied when the submodels have special structural interactions; in this case, we can implicitly store the state space using negligible memory, and the overhead is reduced to $O(K)$, dependent only on the number of submodels, not on the size of their state spaces. On the numerical side,

we have developed algorithms, that exploit the “ultra-sparsity” of the matrices involved in the Kronecker products (they have an average of about one entry per row in practical applications). This has resulted in methods that can outperform the best-known Kronecker-based method (due to W. Stewart and B. Plateau) to compute the vector-matrix products required for the solution of the Markov chain. The results are being implemented in a prototype.

In the short-term future, we plan to confirm our theoretical complexity figures by measuring the performance of our prototype implementation on a suite of realistic applications. Further research will explore distributed algorithms and implementations for Kronecker-based solutions.

This research was conducted in collaboration with Andrew Miner (College of William and Mary), Marco Tilgner (Tokyo Institute of Technology, Japan), Susanna Donatelli (Università di Torino, Italy), and Peter Buchholz and Peter Kemper (Universitaet Dortmund, Germany).

THOMAS W. CROCKETT

Parallel graphics libraries for runtime visualization on massively parallel architectures

Applications which run on large-scale parallel computing systems often generate massive amounts of output data. To assist in the analysis and understanding of this data, we are developing parallel algorithms and software which allow application programs to generate visual output and deliver it across the network to a user’s workstation at runtime.

We have incorporated our research results into a parallel polygon rendering system called PGL. PGL has progressed to the point that an initial version has been released to the high performance computing community. This version includes support for the Intel Paragon and IBM SP2, and is designed to be easily portable to other similar architectures. Documentation and further information are available online at <http://www.icas.edu/~tom/PGL>.

In the near term, we plan to port PGL to workstation networks and more recent parallel architectures such as the Cray T3E. We also plan to add needed features such as transparency, a sphere renderer, and dynamic color quantization. A number of algorithmic issues remain to be explored, including the important issue of scalability to teraflops architectures with thousands of processors.

Portable user interfaces for parallel rendering

Visualization techniques are most helpful when users can direct or interact with the computation or datasets they are viewing. In most parallel rendering applications, the display device is located remotely from the parallel system, necessitating the exchange of image data and control information between the user’s workstation and the supercomputer. The PGL rendering system described above incorporates a simple “proof-of-concept” of such a distributed user interface based on the X11 window system and UNIX networking facilities.

As the power of personal computers begins to rival that of workstations, and as new operating systems such as Windows NT begin to make inroads in the technical community, it is important to provide user interfaces which will work on many different platforms. Sun’s Java language is being touted by many as the best available solution for this problem. In this project we are assessing Java’s utility as the user interface component in a distributed visualization environment. Performance is a critical issue, so our first step is to develop a Java implementation of the data-intensive image display components of PGL’s user interface. To date, we have rewritten the basic

network communication components in Java, and are currently working on the image decompression and display sections of the code.

In the near term, we will conduct performance comparisons between the native X11 implementation and the Java equivalent. If the results are satisfactory, we will then incorporate additional functionality into the user interface to provide an improved graphical environment for monitoring and interacting with parallel applications.

This work was completed in collaboration with Brandon G. Hill.

Fast algorithms for color image quantization

Color quantization refers to the process of reducing the color resolution of a digital image, e.g., to compress the image or to display it on monitors with limited color capabilities. The goal is to find a limited color palette which will serve as an acceptable substitute for the much larger number of colors found in the original image. In earlier work with Shahid Bokhari and David Nicol, we adapted Bokhari's Parametric Binary Dissection algorithm for use as a color-space partitioner, one of the critical steps in most color quantization strategies.

We have recently revisited this problem, with the twin goals of improving both the speed of the partitioner and the quality of the result. By carefully analyzing the problem and incorporating a number of algorithmic improvements, we have been successful on both fronts. Our current implementation provides a 50% improvement in image quality while reducing the partitioning time by 30%, as compared to our original quantizer. We can convert an image with a quarter million pixels from 24-bit to 8-bit color in approximately 0.1 second on a Sun UltraSPARC workstation; the partitioner itself requires only about 10% of the time, with the remainder devoted to image-space operations.

While a fast sequential quantizer is useful in its own right, our ultimate goal is to provide real-time quantization of distributed images to support parallel rendering applications. While the partitioner itself does not appear amenable to parallelization, many of the pixel-level operations are highly parallel. Since these account for the majority of the execution time, we are hopeful that our current strategy will show significant performance improvements in the parallel environment.

STEPHEN GUATTERY

An exact connection between eigenvalues of Laplacian matrices and graph embeddings

Connections between Laplacian spectra (specifically the second smallest eigenvalue λ_2 and its corresponding eigenvector) and properties of the corresponding graph have applications in algorithms. Examples include finding small separators or determining expansion properties; bounds on λ_2 are also useful in analysis of spectral partitioning. Laplacians also often occur in finite difference, finite element, and control volume representations of physical problems involving elliptic partial differential equations. These problems usually include a zero Dirichlet boundary condition that is represented in the Laplacian by deleting the rows and columns corresponding to the boundary vertices. The resulting matrix is positive definite, and its smallest eigenvalue is of interest. A number of techniques using graph embeddings to give lower bounds on the smallest nontrivial eigenvalues of Laplacians have been developed (an example is the path resistance method recently introduced by Steve Guattery (ICASE), Tom Leighton (MIT), and Gary Miller (CMU), which works

for both the graph (positive semidefinite) and zero boundary (positive definite) cases). A common attribute of these techniques is that they do not provide tight lower bounds; we set out to determine why this gap exists, and whether it can be improved.

We have determined that the gap in these bounds is a result of the representation of the problem. Further, we have shown that, by slightly modifying the representation, we can show an exact relationship between the Laplacian eigenvalues and the eigenvalues of a matrix constructed from a specific embedding. This result has an interpretation in terms of resistive circuits and Kirchoff's and Ohm's laws. We have also generalized the result to all invertible real symmetric matrices, though the resistive circuit analogy is lost in that case.

This result opens a number of interesting questions. One line of future research is investigating whether the embedding gives any insight into the structure of the Laplacian's eigenvectors. Another direction is finding easier or faster ways to compute either the eigenvalues or better lower bounds.

This work was done in collaboration Gary Miller of CMU.

Good separator algorithms for weighted graphs

Graph separator algorithms are often used to partition meshes for parallelization. The edges in the resulting separator represent interprocessor communication needs. In many cases, all mesh edges represent similar data dependencies, so the total cost of a partition is the number of edges cut. This can be represented by assigning all edges uniform weights and computing an unweighted separator. However, the problem is sometimes complicated because some edges are involved in different types of computation. Consider the case in which strong connections in a grid are used to determine paths, and which includes a preconditioning step consisting of line solves on these paths. Clearly the cost of interprocessor communication for path edges will be greater than the cost for edges not in the paths. Given an appropriate measure of the relative communication costs, we can apply a weighted separator algorithm in such cases. Our goal is to find a way to assign such costs and generate small weighted separators that give reasonably well balanced partitions.

As a first step in this direction, we have generated partitions for two example grids using METIS, a partitioning program developed by Karypis and Kumar at University of Minnesota. We have varied weights assigned to the path edges and run METIS using various options to determine tradeoffs between the total number of edges cut in forming a partition versus the number of path edges cut. Two things stand out in our initial results: First, METIS's recursive bisection option has given consistently better results than its k-way partitioning option, though at the cost of increased running time of the separator algorithm. Second, the complexity of the multilevel partitioning algorithm can cause it to produce a wide range of results. In particular, the interaction of balance constraints on partition sizes and aspects of the coarsening and uncoarsening cycle can lead to counterintuitive results. For example, increasing the weight of path edges can produce heavier-weight partitions that cut more path edges.

There are many issues left to explore. First, we plan to compare our initial results with the results we get by using other available separator packages such as Chaco. Second, we plan to try and develop a reasonable metric for partition quality that takes into account partition balance, number of edges cut, and number of path edges cut. Third, we plan to examine METIS code and intermediate results to see if it is possible to choose options to get better results more consistently. This may improve our understanding of multilevel separator algorithms by helping us isolate the

effects of different parts of the algorithm.

This is joint work with Alex Pothén (ICASE/Old Dominion University) and Dimitri Mavriplis (ICASE).

Preconditioning for diagonally dominant symmetric matrices

Finding good preconditioners is an important problem in the study of iterative solution methods. In his Ph.D. thesis, Keith Gremban has proposed a technique he calls support tree preconditioners for conjugate gradient applied to diagonally dominant, real symmetric positive definite matrices. Such a preconditioner is constructed on a set of vertices that is a superset of the grid points of the original problem. The vertices of the original problem are the leaves of the support tree; each added vertex v corresponds to the separator between the subgraphs in the original graph induced by the leaves of the subtrees rooted at the children of v . Edge weights in the tree are computed from the sizes of the separators between subgraphs of the original grid induced by subtrees of the support trees. Support trees are designed for parallel execution, since tree-structured matrices can be solved quickly in parallel. Support trees have been shown to work well for some problems, but also have drawbacks that hinder their general acceptance: Construction of the support trees requires recursive bisection of the graph, which can be time-consuming, and the vertex sets they use are larger than that of the original problem, which raises space concerns. Our goal is to reduce these barriers, and to provide improved analysis of their behavior on unstructured grids.

We have worked on applying the techniques from Gremban's Ph.D. thesis, including some embedding techniques for bounding condition numbers of preconditioned systems. We have examined spanning tree preconditioners, which are easy to solve in parallel, but have limitations in the improvements they make in rate of convergence. We have also started studying the analysis of support tree preconditioners.

Several topics for further research remain: The cost of computing separators is high. Can this be reduced through the use of multiway partitioning techniques? Multiway partitioning may also increase the branching of support trees and keep their size down. Also of interest is the relationship of support trees to domain decomposition. In particular, it may be possible to apply Gremban's analysis techniques. Finally, there are a number of open questions concerning the analysis of support tree preconditioning for unstructured grids.

This is joint work with Alex Pothén of ICASE and ODU.

MATTHEW HAINES

Designing efficient substrate software

Substrate software refers to components that are designed to provide a basic set of abstractions upon which other software is built upon. Examples include a filesystem, a message-passing library, and a threads library. It is often expected that substrate software be flexible, efficient, and portable. However, since most substrate software components are designed as "black boxes," they often fail to satisfy the flexibility requirement. The objective of this research is to examine and test alternative design methodologies that improve the flexibility of substrate software without hindering its performance or portability.

Open Implementation Analysis and Design (OIA/D) is a design methodology that attempts to improve the designs of substrate software. To examine this approach, we have used OIA/D to design

and implement a lightweight threads package called OpenThreads. The goals of OpenThreads are improved performance, portability, and flexibility. As a test of its flexibility, we used OpenThreads to implement several platform-independent optimizations for the Orca language runtime system.

We plan to apply this design approach to other substrate components, such as communication libraries, and continue to measure its effectiveness. We also plan to explore alternative methods for improving the design of substrate software.

This research was conducted in collaboration with Koen Langendoen from the Vrije Universiteit in The Netherlands.

VICTORIA INTERRANTE

Visualizing multiple layered surfaces and 3D flow in volume data

Facilitating an intuitive and comprehensive appreciation of the global and local characteristics of flow in a 3D volumetric dataset remains a challenging open problem.

We have developed an efficient new method for simultaneously conveying the 3D shapes and relative depths of multiple layered surfaces in scalar volume data, based on the use of principal direction-driven 3D line integral convolution (LIC), and are continuing to investigate techniques for the more effective portrayal of steady 3D flow via volume LIC.

Future efforts will be directed at improved methods for visualizing 3D unsteady flow, and for facilitating the discrimination of multiple superimposed transparent structures.

This research has been aided by collaboration with Kwan-Liu Ma (ICASE), Chester Grosch (Old Dominion University), and Ron Nowaczyk (Clemson University).

JIM E. JONES

Multigrid convergence acceleration

Multigrid methods are able to solve a widening class of problems with computational work equivalent to a few evaluations of the discrete residual (i.e., work comparable to a few sweeps of a simple relaxation process like Gauss-Seidel or SOR). However, this level of efficiency, sometimes called *textbook multigrid performance*, is typically not reached in the multigrid codes used in complicated fluid dynamics applications. One factor that can cause a loss in multigrid efficiency, is the presence of anisotropies. Here the discrete connections between grid points are much larger in one direction than in others. In fluid dynamics applications, anisotropies can arise because a grid is made finer in one direction in order to resolve a boundary layer. The objective of this work is to develop multigrid solvers which are robust in the presence of anisotropies.

Efficient and robust multigrid solvers for anisotropic problems typically use special implicit smoothers: line relaxation in 2D and plane relaxation in 3D. We have developed structured 3-D multigrid solvers for scalar elliptic problems which use plane relaxation and are robust with respect to anisotropies. For the plane relaxation, we use a 2-D multigrid solver to approximately invert the planes. However, plane relaxation may be difficult to implement in codes using multi-block structured grids where there may be no natural definition of a global ‘plane’. These multi-block structured grids are often used in fluid dynamic applications to capture complex geometries and/or to facilitate parallel processing. We have extended our previous 2-D work to 3-D, showing that

plane relaxation *within the blocks* can be effective provided the size of the blocks and their overlap is proportional to the strength of the anisotropy.

The next step in this research is to test these plane relaxation multigrid ideas in some of the CFD codes at NASA that use multi-block structured grids.

The research was conducted in collaboration with N. Duane Melson (NASA Langley Research Center).

DAVID E. KEYES

Parallelization of implicit CFD codes

The past six months' experience reinforces earlier recommendations of the Newton-Krylov-Schwarz (NKS) framework for the parallelization of implicit computational fluid dynamics codes. The crux of NKS is a balance of convergence rate and data locality, which are fundamentally antithetical in problems with elliptic character (which may arise from either the incompressibility constraint or diffusion). Locality is growing more critical to achieving the performance promised by contemporary machines with pronounced memory hierarchies, in which both vertical (memory/cache) and horizontal (processor/processor) transfer of data and/or data validation information exacts a performance price. Nevertheless, no sequentially optimal algorithms are purely local. Our objective is to design and tune efficient parallel domain decomposition algorithms with inputs from both analysis and architecture.

The matrix-vector products, DAXPYs, and inner products of Newton and Krylov methods are reasonably scalable on today's tightly coupled architectures, but the convergence rates and overall parallel efficiencies of additive Schwarz methods are notoriously dependent upon the number and shape of the subdomains. Except in certain cases for which effective sets of coarse grid operators and intergrid transfer operators are known, so that optimal multilevel preconditioners can be constructed, the number of iterations to convergence generally increases with increasing numbers of subdomains. The communication overhead per iteration increases as well, problem size being held constant. In practical large-scale applications, however, the convergence rate degradation of fine-grain single-level additive Schwarz is often not as serious as the scalar, linear elliptic theory would suggest. Its effects are mitigated by several factors, including an outer context of nonlinearity and pseudo-transient continuation, and strong intercomponent coupling that can be captured exactly in a point-block ILU preconditioner. Another "forgiveness factor" for additive Schwarz, in practice, is the convenience with which Schwarz-based preconditioners can be made to play to the cache in modern microprocessors. As a specific illustration, a three-dimensional unstructured grid Euler code from NASA, parallelized with the PETSc library of Argonne for distributed memory, suffered a factor of 1.6 in iteration count in going from 1 to 32 nodes of the SP2 on a fixed-size problem with ILU(0) as the subdomain preconditioner. However, processor utilization (on a per point, per iteration basis) improved by a factor of 1.4 over this same range, even after taking communication into account, leading to an overall fixed-size wallclock efficiency in excess of 80 percent, which is very respectable for an implicit method. Experimentation on a given processor and a given network leads fairly readily to a subdomain size that optimally balances workingset size and communication-to-computation ratio; this size is the logical unit by which to select the number of subdomains for a given size problem of the same characteristics.

We will continue to work with the PETSc developers to suggest and test new features that serve the needs of computational fluid dynamicists (among others). Our experiences with full potential, Euler, and combustion codes are responsible for many existing and evolving features of the package, which currently runs on more than ten architectures, including ICASE machines and the NASA CAS HPCCP testbeds. In so doing, we will continue to hypothesize about and test programming practices, partly built into the PETSc library and partly dependent upon reorganization of the application code, that permit more of the observed per point, per iteration improvements. We also expect to continue the complementary work of incorporating coefficient information into partitioning and ordering decisions to improve convergence rates. Barry Smith and Lois Curfman McInnes of the PETSc development team and Dinesh Kaushik and Nilan Karunaratne of Old Dominion University are our principal partners in this endeavor.

SCOTT T. LEUTENEGGER

Data base support for subset retrieval and visualization of scientific data

The objective is to design and implement a prototype database to facilitate retrieval of subsets of large scientific data sets. The data subsets are anticipated to be used as inputs to other codes, such as for visualization or MDO.

Currently many scientist store and retrieve data sets as files. When the scientist is interested in a subset of the data they read in the entire data set and strip out the portion of interest. This is not practical when data sets are large. Our approach is to provide database support to retrieve only those pages from disk that contain the desired data. Typical CFD data sets are two or three dimensional, thus we provide multi-attribute indexing techniques.

During a three week visit to ICASE in summer of 96 and one week in December we continued development of our prototype by focusing on support for irregular grids. We revised an earlier prototype to be more efficient in terms of time and space. Next, in collaboration with Kwan-Liu Ma, we interfaced the tool with a multi-resolution interactive visualization tool.

We will continue development of the prototype by completing our revised prototype and conducting an experimental study comparing with octree indexing methods. The comparison study will be conducted in collaboration with ICASE staff scientist Kwan-Liu Ma.

KWAN-LIU MA

A scalable parallel cell-projection volume rendering algorithm for three-dimensional unstructured data

Visualizing three-dimensional unstructured data from aerodynamics calculations is challenging because the associated meshes are typically large in size and irregular in both shape and resolution. The goal of this research is to develop a scalable parallel volume rendering algorithm for the next generation of massively parallel distributed-memory supercomputers which may consist of thousands of very powerful processors.

We use cell-projection instead of ray-casting to provide maximum flexibility in the data distribution and rendering steps. Effective static load balancing is achieved with a round robin distribution of data cells among the processors. A spatial partitioning tree is used to guide the rendering,

optimize the image compositing step, and reduce memory consumption. Communication cost is reduced by buffering messages and by overlapping communication with rendering calculations as much as possible.

By combining the spatial partitioning scheme with techniques which were originally developed for parallel polygon rendering, we have produced a volume renderer for unstructured meshes which employs inexpensive static load balancing to achieve high performance and reasonable efficiency with modest memory consumption. We believe that our algorithm is currently the most effective one available for rendering complex unstructured grids on distributed-memory message-passing architectures. Tests on the IBM SP2 demonstrate that these strategies provide high rendering rates and good scalability. For a dataset containing half a million tetrahedral cells, we achieve 70% parallel efficiency and two frames per second for a 400x400-pixel image using 128 processors. Detailed performance experiments lead us to believe that further improvements are possible.

We plan to conduct additional tests with larger datasets, different image sizes, more processors, and other architectures. In particular, we want to investigate the potential for finer-grained image partitionings and improved termination strategies to enhance the parallel performance of our approach. The ultimate goal is a fast, efficient volume renderer which can handle tens of millions of grid cells using thousands of processors.

This work is being done in collaboration with Tom Crockett.

JOSEPH MANTHEY

Numerical methods for computational aeroacoustics

Numerical schemes for computational aeroacoustics are studied for application to duct acoustics. The primary objective of this work is to develop a time-stable high-order finite-difference scheme suitable for duct acoustic applications. A secondary objective is the time-domain implementation of lined-wall boundary conditions given in the frequency domain.

Many existing high-order finite-difference schemes are not time stable and hence are unsuitable for long-time integration. The primary obstacle to the development of explicit high-order finite-difference schemes is the construction of boundary closures which simultaneously maintain the formal order of accuracy and the numerical stability of the overall scheme. It is proposed that a hybrid seven-point, fourth-order stencil for computing spatial derivatives be used in conjunction with a family of optimized low-dissipation and low-dispersion Runge-Kutta time-marching schemes. An eigenvalue stability analysis has been performed for the hybrid stencil with physical boundary conditions applied. It is found that artificial damping or numerical filtering is necessary for the time stability of the hybrid stencil. The hybrid stencil together with artificial damping or numerical filtering constitute a time-stable high-order finite-difference scheme suitable for duct acoustic applications.

Lined-wall boundary conditions are often formulated in the frequency domain since acoustic response is a function of wave frequency. In order to compute the acoustic modes directly from the Euler equations, however, it is necessary that the impedance boundary condition be reformulated in the time domain. It is proposed that a specific impedance condition, cited frequently in the literature as a model for point-reacting liners, be rewritten in the time domain in a simple algebraic form involving the values of pressure and normal velocity at previous times.

Future work is to investigate the numerical stability of time-domain implementations of lined-wall impedance conditions.

PIYUSH MEHROTRA

A distributed computing environment for ICASE

Distributed heterogeneous computing is being increasingly applied to a variety of large size computational problems. Such computations, for example, the multidisciplinary design optimization of an aircraft, generally consists of multiple heterogeneous modules interacting with each other to solve the problem at hand. Such applications are generally developed by a team in which each discipline is the responsibility of experts in the field. The objective of this project is to develop a GUI based environment which supports the multi-user design of such applications and their execution and monitoring in a heterogeneous environment consisting of a network of workstations, specialized machines and parallel architectures.

The overall goal is to design an interface which is easy to use, easily accessible, and portable. We are planning to leverage off of technologies where available to achieve these goals. In particular, integrating the interface in a web browser, e.g., Netscape, will provide users with a familiar interface on desktops ranging from Unix based workstations, to Windows based PCs, and Macintoshes. It will also allow multiple users to interact in the development and monitoring phases of the process. The use of Java as the programming language for the system will also support heterogeneous portability along with security for access control. The environment will consist of three integrated interfaces supporting the specification of applications, the resource allocation and execution of applications and the monitoring and control of applications. We have completed the design of an initial prototype which focuses on the specification and execution phases and are in the process of implementing the same.

After completing the initial prototype, we will release the environment to selected scientists at ICASE whose feedback will allow us to improve the interface design. We will also continue extending the environment to provide monitoring and control.

This work is being done in collaboration with Kurt Maly, Mohammed Zubair, and graduate students Zikhai Chen and Georgia Liu from Old Dominion University.

Multithreaded system for distributed environments

Traditionally, lightweight threads are supported only within the single address space of a process, or in shared memory environments with multiple processes. Likewise, interprocess communication systems do not currently allow messages to be sent directly entities within a process. The objective of this project is build a system which combines standard interfaces for lightweight threads, pthreads, and interprocess communication, MPI, to support point-to-point communication between any two threads in a distributed memory system.

The Chant runtime system has been built using four layers: point-to-point communication, remote service requests, remote thread operations, and collective communication for thread groups. In the last few months we have been exploring extensions which will support load-balancing of distributed computations via migration of light-weight threads. The load balancing layer consists of several sub-layers. The first is the thread migration layer which supports the actual migration

of a thread from one process to another. The second is the information gathering sub-layer which collects load information as directed by the user along with patterns of communication exhibited by the threads. The last is the decision making layer which controls when, where and which threads to migrate based on the information gathered by the above layer. We have completed the thread migration layer and are now focusing on the design of the other two layers.

After all the layers have been implemented, we plan to use the system for a wide range of applications including a branch and bound algorithm, e.g., the fifteen-puzzle, an adaptive unstructured grid-code and a parallel rendering code. These implementations would allow us to optimize our system.

This work is being done in collaboration with David Cronk, a VILaP graduate student and Matthew Haines from University of Wyoming.

Evaluation and extension of HPF

The goal of High Performance Fortran (HPF) is to provide the user with a high-level language interface for programming scalable parallel architectures and delegating to the compiler the task of producing an explicitly parallel message-passing program. In some applications, this approach may result in dramatic performance losses. An important example is the inspector/executor paradigm, which HPF uses to support irregular data accesses in parallel loops. In many cases, the compiler does not have sufficient information to decide whether an inspector computation is redundant or needs to be repeated. In such cases, the performance of the whole program may be significantly degraded. The objective of this research is to design new user constructs which will allow the compiler to produce more efficient parallel code for the irregular parallel loops.

We have developed an approach to solve this problem through the introduction of constructs allowing explicit manipulation of communication schedules at the HPF language level. Our method is based upon the concept of schedule variables, whose values are communication schedules computed by an inspector or, alternatively, defined by the user by specifying the access pattern to an array with appropriate high-level directives.

These features are currently being implemented in the framework of the Vienna Fortran Compilation System (VFCS). We plan to apply this methodology to a set of important applications, evaluate its performance, and use the results to adjust the functionality of the language extensions accordingly.

This work is being done in cooperation with John Van Rosendale, Kevin Roe, a VILaP graduate student and Hans Zima, Siegfried Benkner and Viera Sipkova from the University of Vienna.

ALEX POTHEN

A parallel sparse indefinite solver

Large, sparse, symmetric indefinite systems of equations occur in computational structural mechanics, electromagnetics, and linear and nonlinear programming. Our goal is to develop parallel algorithms and software for solving sparse, symmetric indefinite linear equations on distributed-memory multiprocessors.

We continued the development of the first parallel solver known to us for sparse indefinite systems of equations. The multifrontal method was used to organize the factorization to make

effective use of cache and memory accesses. New parallel algorithms and dynamic data structures were developed to deal with the irregular computation caused by sparsity and numerical pivoting. An exhaustive pivoting strategy suitable for parallelism was used. MPI was employed for portability. The solver is able to accurately solve indefinite systems from structural analysis and linear programming.

We are tuning our solver for improving the performance, and implementing ordering algorithms suitable for indefinite problems. A complex version for Helmholtz problems is being created. We are collaborating with experts in acoustics and electromagnetics to solve indefinite problems in these application areas.

This is joint work with Florin Dobrian of Old Dominion University, and Yogin Campbell of AT&T.

Parallel Algorithms for Incomplete Factorization Preconditioners

The parallel computation of robust preconditioners is a priority for solution of large systems of equations in unstructured grids and other applications. We are developing algorithms and software that can compute the preconditioners in time proportional to the number of floating point operations and memory accesses. Since these algorithms are parallelizable, implementations on parallel machines will be considered.

We have developed a structure theory based on paths in the adjacency graph of the matrix to predict where zero elements become nonzeros in incomplete factorization (fill elements). A level function is used in incomplete factorization to control the number of fill elements, and we relate the level of fill to lengths of appropriately defined paths in the adjacency graph. This result permits us to search in the neighborhood of a vertex in the graph to predict all fill elements associated with that vertex. We have implemented preliminary versions of these algorithms and have verified that they are more efficient than current incomplete factorization preconditioners in PETSc, the Portable, Extensible Toolkit for Scientific Computing from Argonne National Laboratories. We are also developing new classes of preconditioners based on the elimination tree data structure widely used in direct methods to model the dependencies in the factorization.

This work is at a preliminary stage. We will continue to develop the theory, algorithms, and software for fast computation of preconditioners, and then investigate parallel implementations. We anticipate that this work would lead to faster preconditioning in packages like PETSc.

This is joint work with David Hysom and Gary Kumfert of Old Dominion University.

HANS ZIMA

Integrated compilation environments

The long term commercial success of high performance scalable computing systems will depend on the availability of programming interfaces that combine user-friendliness with efficient target code generation. HPF is a step in the right direction; however, current HPF compilation systems cannot fully exploit the power of such machines and do not provide support for the automatic selection of data distributions; furthermore, most existing software tools are still based on the message passing paradigm.

We study innovative compilation strategies for the automatic translation of Fortran 90 programs (possibly annotated with HPF distribution or alignment directives for selected data structures) to scalable parallel computing systems. The key idea is based on the integration of three major subsystems: a restructuring system, a performance analysis system, and a knowledge-based parallelization support system. Starting point for this research is the existing VFCS restructuring system and a number of performance tools that have been integrated with VFCS.

Future efforts will focus on the integration of performance prediction and scalability analysis. Furthermore, a detailed design for the knowledge base and the related parallelization support system will be carried out.

This work is being performed in cooperation with Mario Pantano and Karin Neuhold at the University of Vienna, and Xian-He Sun of Louisiana State University.

MOHAMMED ZUBAIR

JAVADC: A web-Java based environment to run and monitor parallel applications

Parallel and distributed computing on a cluster of workstations is being increasingly applied to a variety of large size computational problems. Several software systems have been developed that make distributed computing available to an application programmer. Currently we lack an easy-to-use interface through which we can use these software systems from anywhere on the Internet and from any machine (platform independent). The objective of this project is to build a Web-Java based environment which allows users to execute parallel programs on a network of workstations.

In this environment, a user in one Internet domain can configure a parallel environment on a high-performance workstation cluster, HPC, in another domain, run an application on HPC and monitor its progress. We have developed a prototype to support pPVM based applications. The pPVM is a software system which enables distributed application on a set of heterogeneous machines communicating over parallel networks. The Web-Java interface is integrated with the Netscape Web browser to keep the user interface familiar and easy to learn. The use of Java makes the graphical interface portable to different platforms.

In the future, we plan to extend JAVADC to support MPI based applications.

This work was done in collaboration with Kurt Maly, and graduate students Zikhai Chen, and Praveen Vangala from Old Dominion University.

REPORTS AND ABSTRACTS

Özturan, Can: *Worst case complexity of parallel triangular mesh refinement by longest edge bisection*. ICASE Report No. 96-56, (NASA CR-201604), December 30, 1996, 9 pages. To appear in the Proceedings of the 8th SIAM Conference on Parallel Processing for Scientific Computing.

Turkel, E., V.N. Vatsa, and R. Radespiel: *Preconditioning methods for low-speed flows*. ICASE Report No. 96-57, (NASA CR-201605), October 25, 1996, 20 pages. AIAA-96-2460-CP.

Banks, H.T., M.A. Demetriou, and R.C. Smith: *Utilization of coupling effects in compensator design for structural acoustic systems*. ICASE Report No. 96-58, (NASA CR-201608), October 15, 1996, 32 pages. Submitted to the Journal of the Acoustical Society of America.

Jameson, Leland: *Wavelet-based grid generation*. ICASE Report No. 96-59, (NASA CR-201609), October 15, 1996, 18 pages. Submitted to the Journal of Applied and Numerical Mathematics.

Ma, Kwan-Liu, and Victoria Interrante: *Extracting features from 3D unstructured meshes for interactive visualization*. ICASE Report No. 96-60, (NASA CR-201610), November 19, 1996, 13 pages. Submitted to the 1997 Symposium on Interactive 3D Graphics.

Das, Indraneel: *An interior point algorithm for the general nonlinear programming problem with trust region globalization*. ICASE Report No. 96-61, (NASA CR-201615), October 18, 1996, 35 pages. To be submitted to the Journal of Optimization Theory and Applications.

Das, Indraneel, and John Dennis: *Normal-boundary intersection: An alternate method for generating Pareto optimal points in multicriteria optimization problems*. ICASE Report No. 96-62, (NASA CR-201616), December 9, 1996, 34 pages. Submitted to the SIAM Journal of Optimization.

Ma, Kwan-Liu, Scott Leutenegger, and Dimitri Mavriplis: *Interactive exploration of large 3-D unstructured-grid data*. ICASE Report No. 96-63, (NASA CR-201618), October 25, 1996, 18 pages. Submitted to 1997 Symposium on Interactive 3D Graphics.

Pavarino, Luca F.: *Preconditioned mixed spectral element methods for elasticity and Stokes problems*. ICASE Report No. 96-64, (NASA CR-201619), October 25, 1996, 26 pages. Submitted to SIAM Journal of Scientific Computing.

Balakumar, P., and P. Hall: *Optimum suction distribution for transition control*. ICASE Report No. 96-65, (NASA CR-201620), December 31, 1996, 27 pages. Submitted to AIAA Journal.

Ristorcelli, J.R.: *Toward a turbulence constitutive relation for rotating flows*. ICASE Report No. 96-66, (NASA CR-201621), November 19, 1996, 34 pages. Submitted to Theoretical and Computational Fluid Dynamics and to be submitted to the Journal of Geophysical Research.

Shen, Han-Wei, Christopher R. Johnson, and Kwan-Liu Ma: *Visualizing vector fields using line integral convolution and dye advection*. ICASE Report No. 96-67, (NASA CR-201624), December 31, 1996, 16 pages. To appear in the Proceedings of the 1996 Symposium on Volume Visualization.

- Fischer, Paul, and David Gottlieb: *On the optimal number of subdomains for hyperbolic problems on parallel computers*. ICASE Report No. 96-68, (NASA CR-201625), December 12, 1996, 14 pages. To appear in the International Journal of Supercomputing.
- Grosch, C.E., J.M. Seiner, M.Y. Hussaini, and T.L. Jackson: *Numerical simulation of mixing enhancement in a hot supersonic jet*. ICASE Report No. 96-69, (NASA CR-201626), December 31, 1996, 60 pages. To appear in Physics of Fluids.
- Cormen, Thomas H., and David M. Nicol: *Performing out-of-core FFTs on parallel disk systems*. ICASE Report No. 96-70, (NASA CR-201627), December 12, 1996, 15 pages. Submitted to Parallel Computing.
- Lewis, Robert Michael, and Virginia Torczon: *Rank ordering and positive bases in pattern search algorithms*. ICASE Report No. 96-71, (NASA CR-201628), December 12, 1996, 26 pages. Submitted to Mathematical Programming.
- Zijal, Robert, and Gianfranco Ciardo: *Discrete deterministic and stochastic Petri nets*. ICASE Report No. 96-72, (NASA CR-201629), December 31, 1996, 25 pages. To be submitted to the Workshop on Measurement, Modeling, and Evaluation of Computer and Communication Systems.
- Cronk, David, Matthew Haines, and Piyush Mehrotra: *Thread migration in the presence of pointers*. ICASE Report No. 96-73, (NASA CR-201630), December 31, 1996, 13 pages. To appear in the Proceedings of the 30th Hawaii International Conference on System Sciences.
- Chanchio, Kasidit, and Xian-He Sun: *Efficient process migration for parallel processing on non-dedicated networks of workstations*. ICASE Report No. 96-74, (NASA CR-201636), December 31, 1996, 28 pages. To appear in the Proceedings of the International Conference on Parallel Processing '96.
- Crockett, Thomas W.: *Beyond the renderer: Software architecture for parallel graphics and visualization*. ICASE Report No. 96-75, (NASA CR-201637), December 31, 1996, 19 pages. Proceedings of the First Eurographics Workshop in Parallel Graphics and Visualization, Alpha Books, Bristol, UK, September 1996, pp. 1-15.
- Iollo, Angelo, and Manuel D. Salas: *Optimum transonic airfoils based on the Euler equations*. ICASE Report No. 96-76, (NASA CR-201638), December 31, 1996, 34 pages. To be submitted to Computers and Fluids.
- Wilson, Robert V., and Ayodeji O. Demuren: *Numerical simulation of turbulent jets with rectangular cross-section*. ICASE Report No. 97-1, (NASA CR-201642), January 31, 1997, 19 pages. To be submitted to the ASME Journal of Fluids Engineering.
- Somani, Arun K.: *Reliability modeling of structured systems: Exploring symmetry in state-space generation*. ICASE Report No. 97-2, (NASA CR-201643), January 31, 1997, 30 pages. To be submitted to Sigmetrics.

Somani, Arun K., and Kishor S. Trivedi: *Boolean algebraic methods for phased-mission system analysis*. ICASE Report No. 97-3, (NASA CR-201644), March 4, 1997, 21 pages. Submitted to IEEE Transactions on Reliability.

Haines, Matthew: *On designing lightweight threads for substrate software*. ICASE Report No. 97-4, (NASA CR-201645), March 4, 1997, 20 pages. To appear in the Proceedings of USENIX Technical Conference, Anaheim, CA, January 1997.

Ciardo, Gianfranco, and Andrew S. Miner: *Storage alternatives for large structured state spaces*. ICASE Report No. 97-5, (NASA CR-201646), March 4, 1997, 22 pages. Submitted to Performance Tools '97.

Aurovillian, Alok, Hong Zhang, and Malgorzata M. Wiecek: *A bookkeeping strategy for multiple objective linear programs*. ICASE Report No. 97-6, (NASA CR-201647), March 14, 1997, 14 pages. Submitted to the Journal of Computers and Operations Research.

Rubinstein, Robert, and Ye Zhou: *Time correlations and the frequency spectrum of sound radiated by turbulent flows*. ICASE Report No. 97-7, (NASA CR-201648), March 14, 1997, 24 pages. Submitted to Physics of Fluids.

Somani, Arun K., and Allen M. Sansano: *Minimizing overhead in parallel algorithms through overlapping communication/computation*. ICASE Report No. 97-8, (NASA CR-201649), March 14, 1997, 28 pages. Submitted to IEEE Transactions on Parallel and Distributed Computing.

Anderson, W. Kyle, and V. Venkatakrishnan: *Aerodynamic design optimization on unstructured grids with a continuous adjoint formulation*. ICASE Report No. 97-9, (NASA CR-201650), March 14, 1997, 46 pages. To be presented at the 35th AIAA Aerospace Science Meeting, January 6-10, 1997, Reno, NV; also submitted to Computers in Fluids.

Burns, John A., and Belinda B. King: *A note on the mathematical modelling of damped second order systems*. ICASE Report No. 97-10, (NASA CR-201657), March 14, 1997, 12 pages. Submitted to the Journal of Mathematical Systems, Estimation and Control.

Sun, Xian-He, and Yu Zhuang: *A high-order direct solver for Helmholtz equations with Neumann boundary conditions*. ICASE Report No. 97-11, (NASA CR-201658), March 26, 1997, 27 pages. Submitted to the International Conference on Supercomputing.

Lewis, Robert Michael: *A nonlinear programming perspective on sensitivity calculations for systems governed by state equations*. ICASE Report No. 97-12, (NASA CR-201659), March 26, 1997, 37 pages. Submitted to SIAM Review.

Chang, H.-C., D. Gottlieb, M. Marion, and B.W. Sheldon: *Mathematical analysis and optimization of infiltration processes*. ICASE Report No. 97-13, (NASA CR-201660), March 26, 1997, 19 pages. Submitted to the Journal of Computational Physics.

Leutenegger, Scott T., Jeffrey M. Edgington, and Mario A. Lopez: *STR: A simple and efficient algorithm for R-tree packing*. ICASE Report No. 97-14, (NASA CR-201661), March 31, 1997, 31 pages. To appear in the 1997 International Conference on Data Engineering.

Jones, Jim E., and N. Duane Melson: *A note on multi-block relaxation schemes for multigrid solvers*. ICASE Report No. 97-15, (NASA CR-201662), March 31, 1997, 12 pages. To be submitted to the 8th Copper Mountain Conference on Multigrid Methods.

Jones, J.E., Z. Cai, S.F. McCormick, and T.F. Russell: *Control-volume mixed finite element methods*. ICASE Report No. 97-16, (NASA CR-201663), March 31, 1997, 28 pages. Submitted to Computational Geosciences.

Horton, Graham: *On the multilevel solution algorithm for Markov chains*. ICASE Report No. 97-17, (NASA CR-201671), March 31, 1997, 24 pages. Presented at the 1996 Copper Mountain Conference on Iterative Methods; submitted to SISC.

Babin, A., A. Mahalov, B. Nicolaenko, and Y. Zhou: *On the asymptotic regimes and the strongly stratified limit of rotating Boussinesq equations*. ICASE Report No. 97-18, (NASA CR-201672), March 31, 1997, 43 pages. Submitted to Theoretical and Computational Fluid Dynamics.

del Rosario, R.C.H., and R.C. Smith: *LQR control of shell vibrations via piezoceramic actuators*. ICASE Report No. 97-19, (NASA CR-201673), March 31, 1997, 20 pages. To appear in the Proceedings of the 7th International Conference on Control and Estimation of Distributed Parameter Systems.

INTERIM REPORTS

Crockett, Thomas W.: *PGL: A parallel graphics library for distributed memory applications*. ICASE Interim Report No. 29, April 1, 1997.

This report documents the first publicly-distributed version of the PGL software system, Release 1.1. It also breaks new ground as the first hypermedia ICASE report, existing only in electronic form as a collection of HTML documents with associated image and animation files.

PGL is a parallel polygon rendering package designed for use in SPMD-style parallel applications running on distributed memory architectures. It exploits the power of the parallel platform to perform rendering operations in place, enabling applications to generate live visual output without having to move large datasets across the network for subsequent post-processing. PGL includes a core set of 3D graphics routines which provide modelling, rendering, and display operations, and auxilliary components which support higher-level graphics and visualization functions, image transport, and user interaction.

PGL is intended primarily to demonstrate the parallel polygon rendering technology developed at ICASE over the last several years, and to serve as a testbed for further experimentation in this area. As a research code, it lacks the full functionality of commercial graphics packages such as OpenGL or PHIGS. Nonetheless, we have found it to be useful in a variety of parallel applications, and expect that additional capabilities and algorithmic improvements will be incorporated over time.

ICASE COLLOQUIA
October 1, 1996 – March 30, 1997

Name/Affiliation/Title	Date
Srolovitz, David, University of Michigan "Understanding Fracture at the Nano- and Micro- Structural Scales"	October 10
Turner, Leaf, Los Alamos National Laboratory "Incompressible, Inhomogeneous Turbulence Lassoed in a Slab"	October 10
Herring, Jackson, National Center for Atmospheric Research (NCAR) "Recent Results in the Theory of Two Dimensional Turbulence"	October 11
Eyink, Gregory, University of Arizona "Variational Principles for Mean Statistics and PDF Modelling of Turbulence"	October 15
Chu, C.K., Columbia University "Type-Insensitive Solution of Mixed Elliptic-Hyperbolic Equations and Applications to Nonlinear Problems"	October 25
Woodruff, Stephen, Florida State University "Large-Eddy Simulations of a Non-Equilibrium Kolmogorov Flow: Giving Turbulence a Swift Kick"	October 28
Hesthaven, Jan, Brown University "Spectral Multidomain Methods for Problems of Gasdynamics and Electromagnetics in Complex Geometries"	October 31
Lighthill, Sir James, University College, London, England "Recent Advances in Interpreting Hearing Sensitivity"	November 12
Doak, Philip, Institute of Sound and Vibration Research, Southampton, England "Fluctuating Total Enthalpy as the Unique Generalized Acoustic Field"	November 14
Ryzhov, Oleg, Rensselaer Polytechnic Institute "Absolute Instability of a 3D Boundary Layer"	November 15
Zaitsev, Nikolai, Russian Academy of Sciences "Artificial Boundary Conditions for the Wave Equation"	November 21
Radvogin, Yulian, Russian Academy of Sciences "Artificial Boundary Conditions for Acoustics Problems in the Moving Medium"	November 22

Name/Affiliation/Title	Date
Wilson, Robert, Old Dominion University "Numerical Simulation of Turbulent Jets with Rectangular Cross-Section"	December 3
Yusof, Jamaludin, Center for Turbulence Research "DNS and Modeling of the Interaction of Massive Particles with Turbulence"	December 5
Anderson, Kyle, Fluid Mechanics and Acoustics Division, NASA LaRC "Aerodynamic Design Optimization on Unstructured Grids with a Continuous Adjoint Formulation"	January 17
Mavriplis, Dimitri, ICASE "Adaptive Meshing Techniques for Viscous Flow Calculations on Mixed-Element Unstructured Meshes"	January 17
Hege, Hans-Christian, ZIB, Scientific Visualization, Germany "Texture-Based Methods for Vector Field Visualization"	January 22
Ristorcelli, J.R., ICASE "Pseudo-Sound Constitutive Relationship for the Dilatational Covariances in Compressible Turbulence: The Pressure-Dilatation"	January 24
Tsynkov, Semyon V., NRC/Aerodynamic and Acoustic Methods Branch, NASA LaRC "Artificial Boundary Conditions for Infinite-Domain Problems"	January 31
Young, Jennifer, University of Virginia/Polymers and Composite Branch, NASA LaRC "The Challenges Facing Computational Materials Science"	February 20
Michielssen, Eric, University of Illinois at Urbana-Champaign "Novel Fast Integral Equation Based Solvers for Computational Electromagnetics"	March 3
Lee, Dohyung, University of Michigan, Ann Arbor "Local Preconditioning of the Euler and Navier-Stokes Equations"	March 12
Simpson, Timothy, Georgia Institute of Technology "Ranged Sets of Design Requirements and Design Specifications: The Case for Design Capability Indices"	March 18
Sachs, Ekkehard, Universitaet Trier, Germany "Optimization Methods and Applications in Optimal Control"	March 27

OTHER ACTIVITIES

On October 7–9, 1996, ICASE and NASA LaRC co-sponsored the Second Industry Roundtable at the Williamsburg Hospitality House in Williamsburg, VA. The objective of the Roundtable was twofold: 1) to expose ICASE/LaRC scientists to industrial research agendas; and (ii) to acquaint industry with the capabilities and technology available at ICASE/LaRC and academic partners of ICASE. Over 20 sessions in Computer Science, Applied and Numerical Mathematics and Fluid Mechanics were held. There were 173 attendees, and an ICASE interim report will be published.

On November 4–8, 1996 and February 10–14, 1997, ICASE and NASA LaRC co-sponsored a Nonlinear Systems and Control Course at the Pearl Young Theater, NASA Langley Research Center. The course instructor was Professor Alberto Isidori, Universita di Roma La Sapienza. Topics covered were internal structure of nonlinear SISO and MIMO systems, stability and stabilization, and regulation and tracking. There were 40 attendees.

On December 11–13, 1996 ICASE, NASA LaRC, Argonne National Laboratory, and The Cornell Theory Center hosted a Bring-Your-Own Code Workshop on the Portable Parallel Solution of PDE's. This workshop was held at ICASE and was designed for computational engineers and scientists with an interest in distributed computation for large-scale problems in partial differential equations. There were 36 attendees.

ICASE STAFF

I. ADMINISTRATIVE

Manuel D. Salas, Director, M.S., Aeronautics and Astronautics, Polytechnic Institute of Brooklyn, 1970. Fluid Mechanics and Numerical Analysis.

Linda T. Johnson, Office and Financial Administrator

Etta M. Blair, Accounting Supervisor

Barbara A. Cardasis, Administrative Secretary

Shelly M. Johnson, Executive Secretary/Visitor Coordinator

Shannon L. Keeter, Technical Publications Secretary

Rachel A. Lomas, Payroll and Accounting Clerk

Emily N. Todd, Conference Manager

Gwendolyn W. Wesson, Contract Accounting Clerk

Leon M. Clancy, Senior System Manager

Bryan K. Hess, Assistant System Manager

Gregory P. Wheeler, System Operator

II. SCIENCE COUNCIL

Geoffrey Fox, Director, Northeast Parallel Architectural Center, Syracuse University.

Dennis Gannon, Professor, Center for Innovative Computer Applications, Indiana University.

Ashwani Kapila, Professor, Department of Mathematics and Science, Rensselaer Polytechnic Institute.

James P. Kendall, Jet Propulsion Laboratory.

Heinz-Otto Kreiss, Professor, Department of Mathematics, University of California at Los Angeles.

Sanjoy Mitter, Professor of Electrical Engineering, Massachusetts Institute of Technology.

Steven A. Orszag, Professor, Program in Applied & Computational Mathematics, Princeton University.

Paul Rubbert, Unit Chief, Boeing Commercial Airplane Group.

Ahmed Sameh, Department Head of Computer Science, University of Minnesota.

Manuel D. Salas, Director, Institute for Computer Applications in Science and Engineering, NASA Langley Research Center.

III. RESEARCH FELLOWS

Dimitri Mavriplis - Ph.D., Mechanical and Aerospace Engineering, Princeton University, 1988. Applied & Numerical Mathematics [Grid Techniques for Computational Fluid Dynamics]. (February 1997 to August 2001)

Piyush Mehrotra - Ph.D., Computer Science, University of Virginia, 1982. Computer Science [Programming Languages for Multiprocessor Systems]. (January 1991 to September 1999)

IV. SENIOR STAFF SCIENTISTS

Sharath S. Girimaji - Ph.D., Mechanical and Aerospace Engineering, Cornell University, 1990. Fluid Mechanics [Turbulence and Combustion]. (July 1993 to August 1997)

Thomas Jackson - Ph.D., Mathematics, Rensselaer Polytechnic Institute, 1985. Fluid Mechanics. (January 1994 to August 1997)

R. Michael Lewis - Ph.D., Mathematical Sciences, Rice University, 1989. Applied & Numerical Mathematics [Multidisciplinary Design Optimization]. (May 1995 to April 1998)

Josip Lončarić - Ph.D., Applied Mathematics, Harvard University, 1985. Applied & Numerical Mathematics [Multidisciplinary Design Optimization]. (March 1996 to February 1997)

Dimitri Mavriplis - Ph.D., Mechanical and Aerospace Engineering, Princeton University, 1988. Applied & Numerical Mathematics [Grid Techniques for Computational Fluid Dynamics]. (February 1987 to January 1997)

J. Ray Ristorcelli - Ph.D., Mechanical and Aerospace Engineering, Cornell University, 1991. Fluid Mechanics [Turbulence Modeling]. (December 1996 to August 1999)

Robert Rubinstein - Ph.D., Mathematics, Massachusetts Institute of Technology, 1972. Fluid Mechanics [Turbulence Modeling]. (January 1995 to January 1998)

David Sidilkover - Ph.D., Applied Mathematics, The Weizmann Institute of Science, 1989. Applied & Numerical Mathematics [Numerical Analysis and Algorithms]. (November 1994 to August 1997)

V. Venkatakrishnan - Ph.D., Mechanical and Aerospace Engineering, Princeton University, 1987. Applied & Numerical Mathematics [Computational Aerodynamics]. (June 1993 to February 1997)

Ye Zhou - Ph.D., Physics, College of William and Mary, 1987. Fluid Mechanics [Turbulence Modeling]. (October 1992 to September 1997)

V. SCIENTIFIC STAFF

Brian G. Allan - Ph.D., Mechanical Engineering, University of California at Berkeley, 1996. Applied & Numerical Mathematics [Multidisciplinary Design Optimization]. (February 1996 to August 1999)

Eyal Arian - Ph.D., Applied Mathematics, The Weizmann Institute of Science, Israel, 1995. Applied & Numerical Mathematics [Multidisciplinary Design Optimization]. (October 1994 to September 1997)

Stephen Guattery - Ph.D., Computer Science, Carnegie Mellon University, 1995. Computer Science [Parallel Numerical Algorithms, including Partitioning and Mapping]. (September 1995 to August 1997)

M. Ehtesham Hayder - Ph.D., Mechanical and Aerospace Engineering, Princeton University, 1988. Fluid Mechanics [Computational Aeroacoustics]. (September 1995 to September 1997)

Victoria L. Interrante - Ph.D., Computer Science, University of North Carolina at Chapel Hill, 1996. Computer Science [Scientific Visualization]. (March 1996 to March 1998)

Leland M. Jameson - Ph.D., Applied Mathematics, Brown University, 1993. Applied & Numerical Mathematics [Multiresolution Schemes]. (August 1996 to August 1998)

Jim E. Jones - Ph.D., Applied Mathematics, University of Colorado-Boulder, 1995. Computer Science [Parallel Multigrid Methods]. (March 1995 to February 1997)

Li-Shi Luo - Ph.D., Physics, Georgia Institute of Technology, 1993. Computer Science [Parallel Algorithms]. (November 1996 to October 1999)

Kwan-Liu Ma - Ph.D., Computer Science, University of Utah, 1993. Computer Science [Visualization]. (May 1993 to August 1999)

J. Ray Ristorcelli - Ph.D., Mechanical and Aerospace Engineering, Cornell University, 1991. Fluid Mechanics [Turbulence Modeling]. (December 1992 to November 1996)

VI. SENIOR SYSTEMS ANALYST

Thomas W. Crockett - B.S., Mathematics, The College of William & Mary, 1977. Computer Science [System Software for Parallel Computing, Computer Graphics, and Scientific Visualization]. (February 1987 to August 2000)

VII. VISITING SCIENTISTS

Angelo Iollo - Ph.D., Aerospace Engineering, Politecnico di Torino, 1996. Marie Curie Fellow, Project SINUS, INRIA-Sophia Antipolis, France. Applied & Numerical Mathematics [Multidisciplinary Design Optimization and Fluid Dynamic Control]. (October 1996 to February 1997)

Ronald H. Nowaczyk - Ph.D., Experimental Psychology, Miami University, 1977. Professor, Department of Psychology, Clemson University. Applied & Numerical Mathematics [Multidisciplinary Design Optimization]. (January 1997 to September 1997)

John Van Rosendale - Ph.D., Computer Science, University of Illinois, 1980. Program Director for New Technologies, Division of Advanced Scientific Computing, National Science Foundation. Computer Science [Parallel Computing]. (July 1994 to March 1999)

VIII. SENIOR RESEARCH ASSOCIATE

David E. Keyes - Ph.D., Applied Mathematics, Harvard University, 1984. Computer Science [Parallel Numerical Algorithms]

IX. CONSULTANTS

Ivo Babuska - Ph.D., Technical Science, Technical University, Prague, Czechoslovakia, 1951; Mathematics, Academy of Science, Prague, 1956; D.Sc., Mathematics, Academy of Science, Prague, 1960. Robert Trull Chair in Engineering, TICAM, The University of Texas at Austin. Applied & Numerical Mathematics [Finite Element Methods Associated With Structural Engineering]

Ponnampalam Balakumar - Ph.D., Aeronautics and Astronautics, Massachusetts Institute of Technology, 1986. Associate Professor, Department of Aerospace Engineering, Old Dominion University. Fluid Mechanics [Stability and Transition]

David Banks - Ph.D., Computer Science, University of North Carolina, 1993. Assistant Professor, Department of Computer Science, Mississippi State University. Computer Science [Graphics and Visualization]

H. Thomas Banks - Ph.D., Applied Mathematics, Purdue University, 1967. Professor, Department of Mathematics, Center for Research in Scientific Computations, North Carolina State University. Applied & Numerical Mathematics [Control Theory]

Richard W. Barnwell - Ph.D., Engineering Mechanics, Virginia Polytechnic Institute and State University, 1968. Professor, Department of Aerospace and Ocean Engineering, Engineering Science and Mechanics. Virginia Polytechnic Institute and State University. Fluid Mechanics [Turbulence Modeling]

Alvin Bayliss - Ph.D., Mathematics, New York University, 1975. Associate Professor, Technological Institute, Northwestern University. Fluid Mechanics [Numerical Solution of the Equations of Fluid Flow and Acoustics]

John A. Burns - Ph.D., Mathematics, University of Oklahoma, 1973. Professor, Virginia Polytechnic Institute and State University. Applied & Numerical Mathematics [Numerical Methods in Optimal Design and Control]

Barbara M. Chapman - M.S., Mathematics, University of Canterbury, Christchurch, New Zealand, 1985. Director, European Institute for Parallel Computing, University of Vienna. Computer Science [Parallel Language Extensions and Optimizations for Parallel Compilers]

Pao-Liu Chow - Ph.D., Applied Mathematics and Mechanics, Rensselaer Polytechnic Institute, 1967. Professor, Department of Mathematics, Wayne State University. Fluid Mechanics [Aeroacoustics and Noise Control]

Gianfranco Ciardo - Ph.D., Computer Science, Duke University, 1989. Assistant Professor, The College of William & Mary. Computer Science [Reliability Models]

Ayodeji O. Demuren - Ph.D., Mechanical Engineering, Imperial College London, United Kingdom, 1979. Associate Professor, Department of Mechanical Engineering and Mechanics, Old Dominion University. Fluid Mechanics [Numerical Modeling of Turbulent Flows]

Geoffrey Fox - Ph.D., Physics, Cambridge University, 1967. Professor, Department of Computer Science, Syracuse University. Computer Science [Networking]

Dennis B. Gannon - Ph.D., Mathematics, University of California, Davis, 1974. Professor, Department of Computer Science, Indiana University-Bloomington. Computer Science [Investigation of Algorithms and Programming Techniques for Parallel Computers]

Mikhail M. Gilinsky - Ph.D., Aerodynamics, Moscow State University, 1965. Research Professor, Department of Mathematics, Hampton University. Fluid Mechanics [Theoretical and Numerical Research on Nozzle-Jet Flows including Turbulence Models, Acoustic Field Calculation, and Numerical Simulation of Nozzle Flows]

David Gottlieb - Ph.D., Numerical Analysis, Tel-Aviv University, Israel, 1972. Professor, Division of Applied Mathematics, Brown University. Applied & Numerical Mathematics [Boundary Conditions for Hyperbolic Systems]

Chester E. Grosch - Ph.D., Physics and Fluid Dynamics, Stevens Institute of Technology, 1967. Professor, Department of Computer Science and Slover Professor, Department of Oceanography, Old Dominion University. Fluid Mechanics [Acoustics]

Matthew D. Haines - Ph.D., Computer Science, Carnegie Mellon University, 1995. Assistant Professor, Department of Computer Science, University of Wyoming. Computer Science [Parallel Programming Environment and Run Time Systems]

Jan S. Hesthaven - Ph.D., Applied Mathematics/Numerical Analysis, Technical University of Denmark, 1995. NSF Postdoctoral Fellow and Visiting Assistant Professor, Division of Applied Mathematics, Brown University. Applied & Numerical Mathematics [Problems Associated with Computational Electromagnetics]

Gene J.-W. Hou - Ph.D., Computational Mechanics, Design Optimization, University of Iowa, 1983. Associate Professor, Mechanical Engineering Department, Old Dominion University. Applied & Numerical Mathematics [Computational Mechanics Design Optimization]

Fang Q. Hu - Ph.D., Applied Mathematics, Florida State University, 1990. Assistant Professor, Department of Mathematics and Statistics, Old Dominion University. Fluid Mechanics [Instability and Transition]

Ashwani K. Kapila - Ph.D., Theoretical and Applied Mechanics, Cornell University, 1975. Associate Professor, Department of Mathematical Sciences, Rensselaer Polytechnic Institute. Fluid Mechanics [Mathematical Combustion]

Frank Kozusko - Ph.D., Computational and Applied Mathematics, Old Dominion University, 1995. Assistant Professor, Department of Mathematics, Hampton University. Fluid Mechanics [Airfoil Design]

Heinz-Otto Kreiss - Ph.D., Mathematics, Royal Institute of Technology, Sweden, 1960. Professor, Department of Applied Mathematics, California Institute of Technology, Applied & Numerical Mathematics [Numerical Solution of Partial Differential Equations]

David G. Lasseigne - Ph.D., Applied Mathematics, Northwestern University, 1985. Assistant Professor, Department of Mathematics and Statistics, Old Dominion University. Fluid Mechanics [Asymptotic and Numerical Methods for Computational Fluid Dynamics]

Scott T. Leutenegger - Ph.D., Computer Science, University of Wisconsin-Madison, 1990. Assistant Professor, Department of Mathematics and Computer Science, University of Denver. Computer Science [System Software Related to Databases for Scientific Data]

Kurt Maly - Ph.D., Computer Science, Courant Institute, New York University, 1973. Kaufman Professor and Chair, Department of Computer Science, Old Dominion University. Computer Science [High Performance Communication]

James E. Martin - Ph.D., Applied Mathematics, Brown University, 1991. Assistant Professor, Department of Mathematics, Christopher Newport University. Fluid Mechanics [Turbulence and Computation]

Sanjoy K. Mitter - Ph.D., Electrical Engineering, Imperial College of Science & Technology, London, 1965. Professor of Electrical Engineering, Co-Director, Laboratory for Information and Decision Systems, Director, Center for Intelligent Control Systems, Massachusetts Institute of Technology. Fluid Mechanics [Control Theory]

Raj Mittra - Ph.D., Electrical Engineering, University of Toronto, 1957. Professor, Department of Electrical Engineering, Penn State. Computer Science [Computational Electromagnetics]

David M. Nicol - Ph.D., Computer Science, University of Virginia, 1985. Professor, Department of Computer Science, Dartmouth College. Computer Science [Scientific Computing on Scalable parallel Computers]

R.A. Nicolaides - Ph.D., Computer Science, University of London, 1972. Professor, Department of Mathematics, Carnegie Mellon University. Applied & Numerical Mathematics [Numerical Solution of Partial Differential Equations]

Alex Pothén - Ph.D., Applied Mathematics, Cornell University, 1984. Professor, Department of Computer Science, Old Dominion University. Computer Science [Parallel Numerical Algorithms]

Ahmed H. Sameh - Ph.D., Civil Engineering, University of Illinois, 1968. Head, William Norris Chair, and Professor, Department of Computer Science, University of Minnesota. Computer Science [Numerical Algorithms]

Chi-Wang Shu - Ph.D., Mathematics, University of California-Los Angeles, 1986. Associate Professor, Division of Applied Mathematics, Brown University. Fluid Mechanics [Computational Aeroacoustics]

Ralph C. Smith - Ph.D., Mathematics, Montana State University, 1990. Assistant Professor, Department of Mathematics, Iowa State University. Applied & Numerical Mathematics [Optimal Control Techniques for Structural Acoustics Problems]

Shlomo Ta'asan - Ph.D., Applied Mathematics, The Weizmann Institute of Science, 1985. Professor, Department of Mathematics, Carnegie Mellon University. Applied & Numerical Mathematics [Numerical Analysis and Algorithm Development]

Siva Thangam - Ph.D., Mechanical Engineering, Rutgers University, 1980. Professor, Department of Mechanical Engineering, Stevens Institute of Technology. Fluid Mechanics [Turbulence Modeling and Simulation]

Lu Ting - Ph.D., Aeronautics, New York University, 1951. Professor, Courant Institute of Mathematical Sciences, New York University. Fluid Mechanics [Nonlinear Acoustic/Structure Interaction Problems]

Virginia Torczon - Ph.D., Mathematical Sciences, Rice University, 1989. Assistant Professor, Department of Computer Science, The College of William & Mary. Computer Science [Parallel Algorithms for Optimization including Multidisciplinary Design Optimization]

Kishor Trivedi - Ph.D., Computer Science, University of Illinois-Urbana, 1974. Professor, Department of Electrical Engineering, Duke University. Computer Science [Performance and Reliability Modeling Methods, Tools and Applications]

George M. Vahala - Ph.D., Physics, University of Iowa, 1972. Professor, Department of Physics, The College of William & Mary. Fluid Mechanics [Group Renormalization Methods for Turbulence Approximation]

Bram van Leer - Ph.D., Theoretical Astrophysics, Leiden State University, The Netherlands, 1970. Professor, Department of Aerospace Engineering, University of Michigan. Applied & Numerical Mathematics [Computational Fluid Dynamics]

Hans Zima - Ph.D., Mathematics, University of Vienna, Austria, 1964. Professor, Institute for Software Technology and Parallel Systems, University of Vienna, Austria. Computer Science [Compiler Development for Parallel and Distributed Multiprocessors]

Mohammad Zubair - Ph.D., Computer Science, Indian Institute of Technology, Delhi, India, 1987. Professor, Department of Computer Science, Old Dominion University. Computer Science [Performance of Unstructured Flow-Solvers on Multi Processor Machines]

X. GRADUATE STUDENTS

Abdelkader Baggag - Department of Computer Science, The University of Minnesota. (September 1995 to Present)

Zhikai Chen - Department of Computer Science, Old Dominion University. (February 1997 to Present)

David C. Cronk - Department of Computer Science, The College of William & Mary. (August 1993 to Present)

Brandon Hill - Department of Computer Science, Old Dominion University. (December 1996 to Present)

Nilan Karunaratne - Department of Computer Science, Old Dominion University. (August 1995 to December 1996)

Gary Kumfert - Department of Computer Science, Old Dominion University. (January 1997 to Present)

Georgia Liu - Department of Computer Science, Old Dominion University. (February 1997 to Present)

Joe L. Manthey - Department of Mathematics, Old Dominion University. (September 1993 to Present)

Deborah F. Pilkey - Department of Engineering Sciences & Mechanics, Virginia Polytechnic Institute and State University. (October 1995 to Present)

Kevin Roe - Department of Computer Science, The College of William & Mary. (May 1995 to Present)

Praveen Kumar Vangala - Department of Computer Science, Old Dominion University. (November 1996 - Present)

Darren M. Wah - Department of Physics, The College of William & Mary. (February 1997 to Present)

Robert V. Wilson - Department of Mechanical Engineering and Mechanics, Old Dominion University. (October 1992 to December 1996)

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE May 1997		3. REPORT TYPE AND DATES COVERED Contractor Report
4. TITLE AND SUBTITLE Semiannual Report. October 1, 1996 through March 31, 1997			5. FUNDING NUMBERS C NAS1-19480 WU 505-90-52-01	
6. AUTHOR(S)				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Institute for Computer Applications in Science and Engineering Mail Stop 403, NASA Langley Research Center Hampton, VA 23681-0001			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) National Aeronautics and Space Administration Langley Research Center Hampton, VA 23681-0001			10. SPONSORING/MONITORING AGENCY REPORT NUMBER NASA CR-201697	
11. SUPPLEMENTARY NOTES Langley Technical Monitor: Dennis M. Bushnell Final Report				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Unclassified-Unlimited Subject Category 59			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) This report summarizes research conducted at the Institute for Computer Applications in Science and Engineering in applied mathematics, fluid mechanics, and computer science during the period October 1, 1996 through March 31, 1997.				
14. SUBJECT TERMS applied mathematics; multidisciplinary design optimization; fluid mechanics; turbulence; flow control; acoustics; computer science; system software; systems engineering; parallel algorithms			15. NUMBER OF PAGES 60	
			16. PRICE CODE A04	
17. SECURITY CLASSIFICATION OF REPORT Unclassified		18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified		19. SECURITY CLASSIFICATION OF ABSTRACT
				20. LIMITATION OF ABSTRACT